arXiv:1612.00889v1 [cs.DS] 2 Dec 2016

New Frameworks for Offline and Streaming Coreset Constructions

Vladimir Braverman * Dan Feldman† Harry Lang‡

Abstract

Let P be a set (called points), Q be a set (called queries) and a function $f : P \times Q \to [0, \infty)$ (called cost). For an error parameter $\epsilon > 0$, a set $S \subseteq P$ with a *weight function* $w : P \to [0, \infty)$ is an ε-coreset if $\sum_{s \in S} w(s) f(s, q)$ approximates $\sum_{p \in P} f(p, q)$ up to a multiplicative factor of $1 \pm \varepsilon$ for every given query $q \in Q$. Coresets are used to solve fundamental problems in machine learning of streaming and distributed data.

We construct coresets for the k-means clustering of n input points, both in an arbitrary metric space and d-dimensional Euclidean space. For Euclidean space, we present the first coreset whose size is simultaneously independent of both d and n. In particular, this is the first coreset of size $o(n)$ for a stream of n sparse points in a $d \geq n$ dimensional space (e.g. adjacency matrices of graphs). We also provide the first generalizations of such coresets for handling outliers. For arbitrary metric spaces, we improve the dependence on k to $k \log k$ and present a matching lower bound.

For M-estimator clustering (special cases include the well-known k-median and k-means clustering), we introduce a new technique for converting an offline coreset construction to the streaming setting. Our method yields streaming coreset algorithms requiring the storage of $O(S + k \log n)$ points, where S is the size of the offline coreset. In comparison, the previous state-of-the-art was the merge-and-reduce technique that required $O(S \log^{2a+1} n)$ points, where a is the exponent in the offline construction's dependence on ϵ^{-1}. For example, combining our offline and streaming results, we produce a streaming metric k-means coreset algorithm using $O(\epsilon^{-2} k \log k \log n)$ points of storage. The previous state-of-the-art required $O(\epsilon^{-4} k \log k \log^6 n)$ points.

*Department of Computer Science, Johns Hopkins University. This material is based upon work supported in part by the National Science Foundation under Grant No. 1447639, by the Google Faculty Award and by DARPA grant N660001-1-2-4014. Its contents are solely the responsibility of the authors and do not represent the official view of DARPA or the Department of Defense.

†Department of Computer Science, University of Haifa

‡Department of Mathematics, Johns Hopkins University. This material is based upon work supported by the Franco-American Fulbright Commission.

1 Introduction

In the algorithmic field of computer science, we usually have an optimization problem at hand and a state-of-the-art or a straight-forward exhaustive search algorithm that solves it. The challenge is then to suggest a new algorithm with a better running time, storage or other feature. A different and less traditional approach is to use data reduction, which is a compression of the input data in some sense, and to run the (possibly inefficient) existing algorithm on the compressed data. In this case, the problem of solving the problem at hand reduced to the computing a problem-dependent compression such that:

1. an existing algorithm that solves the optimization problem on the original (complete) data, will yield a good approximate solution *for the original data* when applied on the compressed data.

2. The time and space needed for constructing the compression and running the optimization algorithm on the coreset will be better than simply solving the problem on the complete data.

There are many approaches for obtaining such a provable data reduction for different problems and from different fields, such as using uniform sampling, random projections (i.e., the Johnson-Lindenstrauss Lemma), compressed sensing, sketches or PCA.

In this paper we focus on a specific type of a reduced data set, called coreset (or core-set) that was originated in computational geometry, but now applied in other fields such as computer vision and machine learning. Our paper is organized into basic sections: results for maintaining coresets over data streams (Section 3) and results for offline coresets (Sections 4-6). We briefly introduce both of these topics the remainder of this section. Many of our results, along with comparison to prior works, are summarized in Table 1 in Section 2.

In the Appendix A (Section 7) we summarize the merge-and-reduce technique that is used in previous approaches [Che09a, HPM04, HPK07, AMR$^+$12, FL11]. In the Appendix B (Section 8) we provide an alternative framework that generalizes our main result (Theorem 3.1), applying to a wide-array of constructions although giving a weaker bound.

1.1 Streaming Results

In the streaming model of computation, the input arrives sequentially. This differs from the standard model where the algorithm is given free access to the entire input. Given a memory that is linear in the size of the input, these models are evidently equivalent; therefore the goal of a streaming algorithm is to perform the computation using a sublinear amount of memory.

Our stream consists of n elements $p_1, \ldots, p_n$. In the streaming model (or more specifically the insertion-only streaming model, since points that arrive will never be deleted), we attempt to compute our solution using $o(n)$ memory. Sometimes the algorithm will be allowed to pass over the stream multiple times, resulting in another parameter called the number of passes. All of our algorithms use polylog(n) memory and require only a single pass.

Prior to the current work, the merge-and-reduce technique due to Har-Peled and Mazumdar [HPM04] and Bentley and Sax [BS80] was used to maintain a coreset on an insertion-only stream. For a summary of this technique, see Section 7 in the Appendix. In this paper we introduce an alternative technique that reduces the multiplicative overhead from $\log^{2a+1} n$ to $\log n$ (here, a is the offline construction's dependence on $1/\epsilon$). While our method is not as general as merge-and-reduce (it requires that the function in question satisfies more than just the "merge" and "reduce" properties, defined in Section 7), it is general enough to apply to all M-estimators. For the special case of our coreset offline construction for M-estimators (introduced in Section 6), we use a more tailored method that causes this to be $\log n$ additive overhead. Therefore our streaming space complexity matches our offline space complexity, both of which improve upon the state-of-the-art.

The offline coreset construction of [FL11] has the following structure: first, a bicriterion approximation is computed. Second, points are sampled according a distribution that depends only on the distances between points of the input and their assigned bicriterion centers. This suggests a two-pass streaming algorithm (which we later combine into a single pass): in the first pass, construct a bicriterion using an algorithm such as [BMO$^+$11]. In the second pass, sample according to the bicriterion found in the first pass. This provides a two-pass algorithm for a coreset using $O(\epsilon^{-2}k \log k \log n)$-space. Our contribution is showing how these two passes can be combined into a single-pass. Using the algorithm of [BMO$^+$11] to output $O(k \log n)$ centers at any time, we show that this is sufficient to carry out the sampling (originally in the second pass) in parallel without re-reading the stream. Our main lemma (Lemma 3.7) shows that the bicriterion, rather than just providing "central" points to concentrate the sampling, actually can be thought of as a proof of the importance of points for the coreset (technically, a bound on the "sensitivity" that we define at the beginning of Section 3). Moreover, the importance of points is non-increasing as the stream progresses, so we can maintain a sample in the streaming setting without using any additional space.

1.2 Offline Results

The name coreset was suggested by Agarwal, Har-Peled, and Varadarajan in [AHPV04] as a small subset S of points for a given input set P, such that any shape from a given family that covers S will also cover P, after expanding the shape by a factor of $(1 + \varepsilon)$. In particular, the smallest shape that covers S will be a good approximation for the smallest shape that covers P. For approximating different cost functions, e.g. the sum of distances to a given shape, we expect that the total weight of the sample will be similar to the number n of input points. Hence, in their seminal work [HPM04], Har-Peled and Mazumdar used multiplicative weights for each point in S, such that the weighted sum of distances from S to a given shape from the family will approximate its sum of distances from the original data. In [HPM04] each shape in the family was actually a set of k points, and the application was the classic k-means problem.

In this paper, we are given an input set P (called *points*), a family (set) Q of items, called *queries* and a function $f : P \to [0, \infty)$ that is called a *cost* function. A coreset is then a *subset* S of P, that is associated with a non-negative weight function $u : S \to [0, \infty)$ such

that, for every given query $q \in Q$, the sum of original costs $\sum_{p\in P} f(p,q)$ is approximated by the weighted sum $\sum_{p\in S} u(p)f(p,q)$ of costs in S up to a multiplicative factor, i.e.,

$$(1-\varepsilon)\sum_{p\in P} f(p,q) \leq \sum_{p\in S} u(p)f(p,q) \leq (1+\varepsilon)\sum_{p\in P} f(p,q)$$

While our framework is general we demonstrate it on the k-means problem and its variant. There are at least three reasons for this: (i) This is a fundamental problem in both computer science and machine learning, (ii) This is probably the most common clustering technique that used in practice, (iii) Many other clustering and non-clustering problems can be reduced to k-means; e.g. Mixture of Gaussians, Bregman Clustering, or DP-means [FFK11, LBK15, BLK15]. In this context we suggest offline coreset constructions for k-clustering queries, that can be constructed in a streaming fashion using our streaming approach, and are:

- of size linear in k (for $d > \log k$) and arbitrary metric space of dimension d. Current coresets that are subset of the input have size at least cubic in k [LS10]. This is by reducing the total sensitivity to $O(1)$ without introducing negative weights that might be conditioned on the queries as in [FL11].

- of size independent of d for the Euclidean case of k-means (squared distances). This is particular useful for sparse input set of points where $d \geq n$, such as in adjacency matrices of graphs, document-term, or image-object matrices. Recent coreset for sparse k-means of [BF15] is of size exponential in $1/\varepsilon$ and thus turn to $O(n)$ when used when the merge-and-reduce tree (where ε is replaced by $O(\varepsilon/\log(n))$. The result of [FSS13] for k-means is exponential in k/ε and fails with constant probability, so also cannot be used with streaming. Another result of [FSS13] suggests a coreset type set for k-means of size $O(k/\varepsilon)$ but which is based on projections that loss the sparsity of the data. Similar sparsity loss occurs with other projection-type compression methods e.g. in [CEM+15]. Nevertheless, we use the technique in [FSS13] to bound the dimension of the k-means problem by $O(k/\varepsilon)$.

- of size independent of d for the Euclidean case, and non-squared distances, using weak coresets. These coresets can be used to approximates the optimal solution, but not every set of k centers. Unlike the weak coresets in [FL11, FMS07], we can use any existing heuristic on these coresets, as explained in Section 6.1.

- Robust to outliers. This is since the general pseudo-metric definition we used (inspired by [FS12], support m-estimators which is a tool for handling outliers [Tyl87]. Unlike in [FS12] our coresets are linear (and not exponential) in k, and also independent of n (and not logarithmic in n).

2 Related Work

The following table summarizes previous work along with our current results. By far, the most widely-studied problems in this class have been the k-median and k-means func-

tions. In general, the extension to arbitrary M-estimators is non-trivial; the first such result was [FS12]. Our approach naturally lends itself to this extension. M-estimators are highly important for noisy data or data with outliers. As one example, Huber's estimator is widely used in the statistics community [HHR11, Hub81]. It was written that "this estimator is so satisfactory that it has been recommended for almost all situations" [Zha11]. Our results work not only for Huber's estimator but for all M-estimators, such as the Cauchy and Tukey biweight functions which are also well-used functions.

Note that in the below table, $\tilde{O}$ notation is used to write in terms of d, ϵ, k, and $\log n$ (therefore hiding factors of $\log\log n$ but not $\log n$).

Problem	Offline Size	Streaming Size	Paper
Euclidean k-means	$O(k\epsilon^{-d}\log n)$	$O(k\epsilon^{-d}\log^{d+2} n)$	[HPM04]
Euclidean k-means	$O(k^3\epsilon^{-(d+1)})$	$O(k^3\epsilon^{-(d+1)}\log^{d+2} n)$	[HPK07]
Euclidean k-means	$O(dk^2\epsilon^{-2}\log n)$	$O(dk^2\epsilon^{-2}\log^8 n)$	[Che09a]
Euclidean k-means	$O(dk\log k\epsilon^{-4})$	$O(dk\log k\epsilon^{-4}\log^5 n)$	[FL11]
Euclidean k-means	$\tilde{O}((d/\epsilon)^{O(d)}k\log n)$	$\tilde{O}((d/\epsilon)^{O(d)}k\log^{O(d)} n)$	[AMR+12]
Euclidean k-means	$O(\epsilon^{-2}k\log k\min(k/\epsilon,d))$	$O(\epsilon^{-2}k\log k\min(\frac{k}{\epsilon},d)+k\log n)$	**
Metric k-means	$O(\epsilon^{-2}k^2\log^2 n)$	$O(\epsilon^{-2}k^2\log^8 n)$	[Che09b]
Metric k-means	$O(\epsilon^{-4}k\log k\log n)$	$O(\epsilon^{-4}k\log k\log^6 n)$	[FL11]
Metric k-means	$O(\epsilon^{-2}k\log k\log n)$	$O(\epsilon^{-2}k\log k\log n)$	**
Euclidean k-median	$\tilde{O}(dk^2\epsilon^{-2}\log n)$	$O(dk^2\epsilon^{-2}\log^8 n)$	[Che09a]
Euclidean k-median	$O(k\epsilon^{-d}\log n)$	$O(k\epsilon^{-d}\log^{d+2} n)$	[HPM04]
Euclidean k-median	$O(k^2\epsilon^{-d})$	$O(k^2\epsilon^{-(d)}\log^{d+1} n)$	[HPK07]
Euclidean k-median	$O(d\epsilon^{-2}k\log k)$	$O(d\epsilon^{-2}k\log k\log^3 n)$	[FL11]
Euclidean k-median	$O(d\epsilon^{-2}k\log k)$	$O(d\epsilon^{-2}k\log k+k\log n)$	**
Metric k-median	$O(k^2\epsilon^{-2}\log^2 n)$	$O(k^2\epsilon^{-2}\log^8 n)$	[Che09a]
Metric k-median	$O(\epsilon^{-2}k\log k\log n)$	$O(\epsilon^{-2}k\log k\log^4 n)$	[FL11]
Metric k-median	$O(\epsilon^{-2}k\log k\log n)$	$O(\epsilon^{-2}k\log k\log n)$	**
Euclidean M-estimator	$O(\epsilon^{-2}k^{O(k)}d^2\log^2 n)$	$O(\epsilon^{-2}k^{O(k)}d^2\log^5 n)$	[FS12]
Euclidean M-estimator	$O(d\epsilon^{-2}k\log k)$	$O(d\epsilon^{-2}k\log k+k\log n)$	**
Metric M-estimator	$O(\epsilon^{-2}k^{O(k)}\log^4 n)$	$O(\epsilon^{-2}k^{O(k)}\log^7 n)$	[FS12]
Metric M-estimator	$O(\epsilon^{-2}k\log k\log n)$	$O(\epsilon^{-2}k\log k\log n)$	**

Table 1: Summary of Related Work

Framework. A generic framework for coreset construction was suggested in [FL11]. The main technique is a reduction from coreset to ε-approximations, that can be computed using non-uniform sampling. The distribution of the sampling is based on the importance of each point (in some well defined sense), and the size of the coreset depends on the sum of these importance levels. This term of importance appeared in the literature as leverage score (in the context of low-rank approximation, see [PKB14] and references therein), or, for the case of k-clustering, sensitivity [LS10]. The proof of many previous coreset constructions were

significantly simplified by using this framework, and maybe more importantly, the size of these coresets was sometimes significantly reduced; see [FL11] for references. Many of these coresets size can be further improve by our improved framework, as explained below.

The size of the coreset in [FL11] depends quadratically on the sum of sensitivities, called *total sensitivity* [LS10]. In this paper, we reduce the size of the coreset that are constructed by this framework to be only near-linear ($t \log t$) in the total sensitivity t. In addition, we generalize and significantly simplify the notation and results from this framework.

k-means. In the k-means problem we wish to compute a set k of centers (points) in some metric space, such that the sum of squared distances to the input points is minimized, where each input point is assigned to its nearest center. The corresponding coreset is a positively weighted subset of points that approximates this cost to every given set of k centers. First deterministic coresets of size exponential in d were first suggested by Har-Peled and Mazumdar in [HPM04]. The first coreset construction of size polynomial in d was suggested by Ke-Chen in [Che09a] using several sets of uniform sampling.

The state-of-the-art is the result of Schulman and Langberg [LS10] who suggested a coreset of size $O(d^2k^3/\varepsilon^2)$ for k-means in the Euclidean case based on non-uniform sampling. The distribution is similar to the distribution in [FMS07] over the input points, however in [FMS07] the goal was to have weaker version coresets that can be used to solve the optimal solution, but are of size independent of d.

Some kind of coreset for k-means of size near-linear in k was suggested in [FL11]. However, unlike the definition of this paper, the multiplicative weights of some of the points in this coreset were (i) negative, and (ii) depends on the query, i.e., instead of a weight $w(p) > 0$ for an input point p, as in this paper, the weight is $w(p, C) \in \mathbb{R}$ where C is the set of queries. While exhaustive search was suggested to compute a PTAS for the coreset, it is not clear how to compute existing algorithms or heuristics (what is actually done in practice) on such a coreset. On the contrary, generalizing an existing approximation algorithm for the k-means problem to handle positively weights is easy, and public implementations are not hard to find (e.g. in Matlab and Python).

k-mean for handling outliers. Coresets for k-means and its variants that handle outliers via m-estimators were suggested in [FS12], and also inspired our paper. The size of these coresets is exponential in k and also depend on logn. For comparison, we suggest similar coreset of size near-linear in k, and independent of n. PTAS for handling exactly m-outliers was suggested in [Che08] but with no coreset or streaming version.

Streaming. The metric results of [Che09a, FL11] and Euclidean results of [Che09a, HPM04, HPK07, FL11] that rely on merge-and-reduce have already been mentioned. A summary of these results appears in the tables below. For the specific case of Euclidean space, a more diverse set of stronger results is known. In particular, coreset constructions are known that do not begin with a bicriterion solution, and whose streaming variant does not rely on merge-and-reduce [AMR$^+$12]. With the additional assumption in Euclidean

space that the points lie on a discrete grid $\{1, \ldots, \Delta\}^d$, alternative techniques are known for k-means and other problems, even when the stream allows the deletion of points [FS05].

3 Streaming Algorithm

We present a streaming algorithm for constructing a coreset for metric k-means clustering that requires the storage of $O(\epsilon^{-2}k \log n)$ points. The previous state-of-the-art [FL11] required the storage of $O(\epsilon^{-4}k \log k \log^6 n)$ points. In this section we assume the correctness of our offline algorithm, which is proven in Section 6.

More generally, our technique works for M-estimators, a general class of clustering objectives that includes the well-known k-median and k-means functions as special cases. Other special cases include the Cauchy functions, the Tukey functions, and the L_p norms. Our method combines a streaming bicriterion algorithm [BMO$^+$11] and a batch coreset construction [FL11] to create a streaming coreset algorithm. The space requirements are combined addivitely, therefore ensuring no overhead.

The streaming algorithm of [BMO$^+$11] provides a bicriterion solution using $O(k \log n)$ space. Our new offline construction of Section 6 requires $O(\epsilon^{-2}k \log k \log n)$ space. Therefore our main theorem yields a streaming algorithm that combines these spaces additively, therefore requiring $O(\epsilon^{-2}k \log k \log n)$ space while maintaining a coreset for k-means clustering. The previous state-of-the-art framework that works for the metric variant (other methods are known to improve upon this for the special case of Euclidean space) was the merge-and-reduce technique [BS80] that yields a streaming algorithm requiring $O(\epsilon^{-4}k \log k \log^6 n)$ space, incurring an overhead of $\Theta(\log^5 n)$ over the offline coreset size. In comparison, our framework incurs no overhead. The additional improvement in our space is due the improved offline construction given in Sections 5-7.

We now state our main theorem. The result is stated in full generality: the k-median clustering in a ρ-metric space (see Definition 6.1). Note that metric k-means clustering corresponds to setting $\rho = 2$. Also, the probability of success $1 - \delta$ typically has one of two meanings: that the construction succeeds at the end of the stream (a weaker result), or that the construction succeeds at every intermediate point of the stream (a stronger result). Our theorem gives the stronger result, maintaining a valid coreset at every point of the stream.

Theorem 3.1 (Main Theorem)**.** *There exists an insertion-only streaming algorithm that maintains a (k, ϵ)-coreset for k-median clustering in a ρ-metric space, requires the storage of $O(\rho^2\epsilon^{-2}k \log(\rho k) \log(n) \log(1/\delta))$ points, has poly$(k, \log n, \rho, \epsilon, \log(1/\delta))$ worst-case update time, and succeeds at every point of the stream with probability $1 - \delta$.*

Our method can be applied to the coreset constructions of [HPM04, HPK07, Che09a, FL11] with a multiplicative overhead of $O(\log n)$. Our second theorem is a more generally applicable technique; it applies to all constructions that first compute a bicriterion solution and then sample points according to the bicriterion solution. The constructions of [HPM04, HPK07, Che09a, FL11] follow this outline, and we are unaware of any constructions which do

not. The theorem yields immediate corollaries as well as reducing certain streaming coreset problems to that of constructing an offline coreset.

Theorem 3.2. *Given an offline algorithm that constructs a (k, ϵ)-coreset consisting of $S = S(n, k, \epsilon, \delta)$ points with probability $1 - \delta$ by sampling points based on a bicriterion solution, there exists a streaming algorithm requiring the storage of $O(S \log n)$ points that maintains a (k, ϵ)-coreset on an insertion-only stream.*

Proof Sketch. The known offline coreset constructions start with a bicriterion solution of $O(k)$ points. We modify the algorithm of [BMO+11] to output $O(k \log n)$ centers; this is trivial since the final step of the algorithm of [BMO+11] is to take the $O(k \log n)$ centers stored in memory and reduce them to exactly k centers to provide a solution. Our first modification to the original algorithm is thus to simply remove this final step, but we must also keep a datastructure storing $\log(1/\epsilon)$ intermediate states of these $O(k \log n)$ centers. See Section 8 for a precise description of our modification and the sampling method, applied to the construction of [FL11] as an example (but equally applicable to [HPM04, HPK07, Che09a]). As the high-level idea, since the bicriterion given to the offline construction consists of $O(k \log n)$ centers instead of exactly k, the number of additional points taken in the coreset increases by a factor of $O(\log n)$. □

Two important corollaries include:

1. Using the result of [FL11], we obtain a streaming algorithm that maintains a (k, ϵ)-coreset with negative weights for metric k-median requiring the storage of $O(\epsilon^{-2} k \log n)$ points.

2. Given a $O(k \cdot \text{poly}(\epsilon, \log n, \log(1/\delta)))$ point (k, ϵ)-coreset, we would obtain a streaming algorithm that maintains a (k, ϵ)-coreset (with only positive weights) for metric k-median requiring the storage of $O(k \cdot \text{poly}(\epsilon, \log n, \log(1/\delta)))$ points. This differs from Theorem 3.1 in that the dependence on k is linear instead of $O(k \log k)$.

3.1 Definitions

We begin by defining a ρ-metric space, which is defined in full as Definition 6.1. Briefly, let X be a set. If $D : X \times X \to [0, \infty)$ is a symmetric function such that for every $x, z \in X$ we have that $D(x, z) \leq \rho(D(x, y) + D(y, z))$ for every $y \in X$, when we call (X, D) a ρ-metric space. Note that this is a weakening of the triangle inequality, and at $\rho = 1$ we recover the definition of a metric space. All M-estimators can be re-cast for a certain constant value of ρ, and k-means is obtained with $\rho = 2$. This generality is therefore useful and working in this language allows us to naturally generalize our results to any M-estimator.

The k-median problem is, given an input set P and an integer $k \geq 1$, to find a set C of k points that minimizes:

$$\sum_{p \in P} \min_{c \in C} D(p, c)$$

We use $\mathtt{OPT}_k(P)$ to denote this minimal value. As this is NP-Hard to compute, we settle for an approximation. The notion of a bicriterion approximation is well-known; we state a definition that suits our needs while also fitting into the definition of previous works.

Definition 3.3 ((α, β)-approximation)**.** *An (α, β)-approximation for the k-median clustering of a multiset P is a map $\pi : P \to B$ for some set B such that $\sum_{p \in P} w(p) D(p, \pi(p)) \leq \alpha \mathsf{OPT}_k(P)$ and $|B| \leq \beta k$.*

We now define a coreset:

Definition 3.4 ((k, ϵ)-coreset)**.** *A (k, ϵ)-coreset of a multiset P is a weighted set (S, v) with non-negative weight function v such that for every $Z \in \mathcal{X}^k$ we have $(1 - \epsilon) \sum_{p \in P} D(p, Z) \leq \sum_{s \in S} v(s) D(s, Z) \leq (1 + \epsilon) \sum_{p \in P} D(p, Z)$.*

Coresets with arbitrary weight functions (i.e. with negative weights allowed) have been considered [FL11, etc]. However, computing approximate solutions on these coresets in polynomial-time remains a challenge, so we restrict our definition to non-negative weight functions. This ensures that an approximate solution can be quickly produced. This implies a PTAS for Euclidean space and a polynomial-time $2\gamma(1 + \epsilon)$-approximation for general metric spaces (where γ is the best polynomial-time approximation factor for the problem in the batch setting). This factor of $2\gamma(1 + \epsilon)$ is well-known in the literature, see [COP03, BMO+11, GMM+03] for details.

3.2 Constant-Approximation Algorithm

Let P_i denote the prefix of the stream $\{p_1, \ldots, p_i\}$. The entire stream is then P_n. Consider the moment when the first i points have arrived, meaning that the prefix P_i is the current set of arrived points. The algorithm $\mathcal{A}$ of [BMO+11] provides an $(O(1), O(\log n))$-approximation of P_i in the following sense. Define $f_0 : \emptyset \to \emptyset$ as the null map, and define $B_i = \text{image}(f_i)$. Upon receiving point p_i, algorithm $\mathcal{A}$ defines a map $f_i : B_{i-1} \cup \{p_i\} \to B_i$. We define $\pi_i : P_i \to B_i$ by $\pi_i(p_j) = f_i(f_{i-1}(\ldots(f_j(p_j))\ldots))$ for each $1 \leq j \leq i$. These mappings have an essential gaurantee stated in the following lemma.

Theorem 3.5 ([BMO+11])**.** *For every $1 \leq i \leq n$, after receiving P_i, Algorithm $\mathcal{A}(k, n, \delta)$ defines a function f_i such that with probability $1 - \delta$, using the above definition of π_i, the bound $\sum_{p \in P_i} D(p, \pi_i(p)) \leq \alpha \mathsf{OPT}_k(P_i)$ holds. The algorithm deterministically requires the storage of $O(k(\log n + \log(1/\delta)))$ points.*

3.3 Offline Coreset Construction

We briefly describe the offline coreset construction. The proof of correctness can be found in Sections 5 and 6. It is this construction that we will maintain in the streaming setting.

The sensitivity of a point $p \in P$ is defined as:

$$s(p) = \max_{Z \in \mathcal{X}^k} \frac{D(p, Z)}{\sum_{q \in P} D(q, Z)}$$

Notice that $0 \leq s(p) \leq 1$. We give an upper bound $s'(p) \in [s(p), 1]$. Define the total sensitivity $t = \sum_{p \in P} s(p)$. Likewise, we give an upper bound $t' \geq t$ where $t' = \sum_{p \in P} s'(p)$ and will show that $t = O(k)$. The sampling probability distribution at point p is set to $s'(p)/t'$. We take an i.i.d. sample from P of size m for any $m \geq ct'\epsilon^{-2}(\log n \log t' + \log(1/\delta))$ where c is a constant.

Let R be the union of these m i.i.d. samples, and then define a weight function $v : R \to [0, \infty)$ where $v(r) = (|R|s'(r))^{-1}$. It is proven as one of our main theorems (Theorem 6.6) that the weighted set (R, v) is a (k, ϵ)-coreset for P.

3.4 Bounding the Sensitivity

Consider the prefix P_i which is the input after the first i points have arrived. Using Algorithm $\mathcal{A}$ we obtain an (α, β)-approximation π_i where $\alpha = O(1)$ and $\beta = O(\log n)$. Recall that B_i is the image of this approximation, i.e. $B_i = \text{image}(\pi_i(P_i))$.

Running an offline (γ, λ)-approximation algorithm on B_i, we obtain a multiset C_i of at most λk distinct points. Let p' denote the element of C_i nearest to $\pi_i(p)$ (this is the element of C_i that p gets mapped to when we pass from $P_i \to B_i \to C_i$). The following lemma implies that $\sum_{p \in P} w(p)D(p, p') \leq \bar{\alpha}\mathsf{OPT}(P)$ where $\bar{\alpha} = \rho\alpha + 2\rho^2\gamma(\alpha + 1)$. This is an observation used widely in the literature [COP03, BMO+11], but we include a proof for completeness.

Lemma 3.6. *Let B be a (α, β)-approximation of P, and let C be a (γ, λ)-approximation of B. Then C is a $(\rho\alpha + 2\rho^2\gamma(\alpha + 1), \lambda)$-approximation of A.*

Proof. Let $\pi : P \to B$ be the (α, β)-approximation of P and let $t : B \to C$ be the (γ, λ)-approximation of B. In the following, all sums will be taken over all $p \in P$. The hypotheses state that $\sum D(p, \pi(p)) \leq \alpha\mathsf{OPT}(P)$ and $\sum D(\pi(p), t(\pi(p))) \leq \gamma\mathsf{OPT}(B)$. Let P^* be an optimal clustering of P, that is $\sum D(p, P^*) = \mathsf{OPT}(P)$. Then $\frac{1}{2}\mathsf{OPT}(B) \leq \sum D(\pi(p), P^*) \leq \rho\sum(D(\pi(p), p) + D(p, P^*)) \leq \rho(\alpha + 1)\mathsf{OPT}(P)$. The factor of $\frac{1}{2}$ comes from the fact that $\mathsf{OPT}(B)$ is defined using centers restricted to B (see [GMM+03] for details). We now write $\sum D(p, t(\pi(p))) \leq \rho\sum(D(p, \pi(p)) + D(\pi(p), t(\pi(p)))) \leq (\rho\alpha + 2\rho^2\gamma(\alpha+1))\mathsf{OPT}(P)$ as desired. □

We now prove the following lemma which gives us our sampling probability $s'(p)$. Recall that for the construction to succeed, the sampling probability $s'(p)$ must be at least the sensitivity $s(p)$ (defined in the previous subsection). Since we focus on a single iteration, we drop subscripts and write $C = C_i$ and $P = P_i$. Let $p \mapsto p'$ be an $(\bar{\alpha}, \lambda)$-approximation of P. Define $P(p) = \{q \in P : q' = p'\}$ to be the cluster containing p.

Lemma 3.7. *Let the map $p \mapsto p'$ define an $(\bar{\alpha}, \lambda)$-approximation for the k-median clustering of P. For every point $p \in P$:*

$$s(p) \leq \frac{\rho\bar{\alpha}D(p, p')}{\sum_{q \in P} D(q, q')} + \frac{\rho^2(\bar{\alpha} + 1)}{|P(p)|}$$

Proof. For an arbitrary $Z \in \mathcal{X}^k$ we need to provide a uniform bound for

$$\begin{aligned}\frac{D(p,Z)}{\sum_{q\in P} D(q,Z)} &\leq \frac{\rho D(p,p')}{\sum_{q\in P} D(q,Z)} + \frac{\rho D(p',Z)}{\sum_{q\in P} D(q,Z)} \\ &\leq \frac{\bar{\alpha}\rho D(p,p')}{\sum_{q\in P} D(q,q')} + \frac{\rho D(p',Z)}{\sum_{q\in P} D(q,Z)}\end{aligned} \tag{1}$$

where the second inequality holds because $\sum_{q\in P} D(q,q') \leq \bar{\alpha}\mathsf{OPT}(P) \leq \sum_{q\in P} D(q,Z)$. To bound the last term, recall that $q' = p'$ for all $q \in P(p)$ so:

$$\begin{aligned}D(p',Z)|P(p)| &= \sum_{q\in P(p)} D(p',Z) = \sum_{q\in P(p)} D(q',Z) \\ &\leq \rho \sum_{q\in P(p)} (D(q',q) + D(q,Z)) \\ &\leq \rho \sum_{q\in P} D(q',q) + \rho \sum_{q\in P(p)} D(q,Z) \\ &\leq \rho\bar{\alpha} \sum_{q\in P} D(q,Z) + \rho \sum_{q\in P(p)} D(q,Z) \\ &\leq \rho(\bar{\alpha}+1) \sum_{q\in P} D(q,Z)\end{aligned}$$

Dividing by $|P(p)| \sum_{q\in P} D(q,Z)$ gives

$$\frac{D(p',Z)}{\sum_{q\in P} D(q,Z)} \leq \frac{\rho(\bar{\alpha}+1)}{|P(p)|}$$

Substituting this in (16) yields the desired result. □

We therefore define our upper bound $s'(p)$ as in the lemma. An immediate but extremely important consequence of Lemma 3.7 is that $t' = \sum_{p\in P} s'(p) = \rho\bar{\alpha} + \rho^2(\bar{\alpha}+1)k \leq 3\rho^2\bar{\alpha}k$. This can be seen by directly summing the formula given in the lemma.

3.5 Streaming Algorithm

We now state Algorithm 1, which we then prove maintains a coreset. To use Lemma 3.7 to determine $s'(p)$, we will compute the cluster sizes $|P(p)|$ and estimate the clustering cost $\sum_{q\in P} D(q,q')$. We must bound the clustering cost from below because we need an upper-bound of $s(p)$. On Line 9, L is an estimate of the cost of clustering P to the centers C. On Line 4, c is the absolute constant used in Theorem 6.6.

Algorithm 1 outputs (R, v) such that point p is sampled with probability $xs'(p)$ where x is defined on Line 4. For each p that has arrived, the value of $s'(p)$ is non-increasing (notice that it is defined as the minimum of itself and a new value on Line 13), so it is possible to

Algorithm 1: Input: stream of n points in a ρ-metric space, $\epsilon > 0$, $k \in \mathbb{N}$, maximum failure probability $\delta > 0$

Initilization:
$R \leftarrow \emptyset$
$t' \leftarrow \rho\bar{\alpha} + \rho^2(\bar{\alpha}+1)k$
$x \leftarrow 2c\epsilon^{-2}(\log n \log t' + \log(1/\delta))$
Initialize $\mathcal{A}(k, n, \delta)$
Update Process: after receiving p_i
$(B_i, f_i) \leftarrow \mathcal{A}.update(p_i)$
$C \leftarrow$ an (γ, λ)-approximation of B_i
$L \leftarrow D(p_i, C) + \bar{\alpha}^{-1}(1+\epsilon)^{-1}\sum_{r\in R} v(r)D(r, C)$
for $r \in R$ **do**
    $\pi_i(s) \leftarrow f_i(\pi_{i-1}(r))$
    $z(r) \leftarrow \frac{\rho\bar{\alpha}D(p,p')}{L} + \frac{\rho^2(\bar{\alpha}+1)}{|P(p)|}$
    $s'(r) \leftarrow \min(s'(r), z(r))$
for $r \in R$ **do**
    if $u(r) > xs'(r)$ **then**
        Delete r from R
$u(p_i) \leftarrow$ uniform random from $[0, 1)$
if $u(p_i) \leq xs'(p_i)$ **then**
    Add p_i to R
    $\pi_i(p_i) \leftarrow f_i(p_i)$
Query Process:
for *each* $r \in R$ **do**
    $v(r) \leftarrow (|R|s'(r))^{-1}$

maintain this in the streaming setting since once the deletion condition on Line 15 becomes satisfied, it remains satisfied forever.

We now proceed with the proof of Theorem 1. Since the probability of storing point p is $xs'(p)$, the expected space of Algorithm 1 is xt'. By Lemma 3.7 that implies $t' \leq 3\rho^2\bar{\alpha}k$, we then bound the expected space as $2c\epsilon^{-2}(\log n \log t + \log(1/\delta))(3\rho^2\bar{\alpha}k)$. Simplifying notation by defining an absolute constant $\tilde{c}$ (a function of c and $\bar{\alpha}$), we write this expected space as $\tilde{c}\rho^2\epsilon^{-2}k\log(\rho k)\log n\log(1/\delta)$. By a Chernoff bound, the high-probability gaurantee follows by replacing $\tilde{c}$ with $2\tilde{c}$.

Proof of Theorem 1. The proof of Theorem 1 can be divided into the following pieces:

1. For correctness (to satisfy the bound given in Lemma 3.7, we must show that $L \leq \sum_{p\in P_i} D(p,p')$. For space, it is important that L is not too small. In particular, the space grows as $1/L$. We show that L is a ϵ-approximation of the true cost.

2. The value of $|P(p)|$ can be computed exactly for every p. This is needed on Line 12.

3. The construction of Algorithm 1 that samples p with probability $xs'(p)$ can be processed to be identical to the offline construction of Subsection 3.3 that takes an i.i.d. sample of size xt' from the distribution where p is given sampling probability $s'(p)/t'$.

1: To lower bound the clustering cost, inductively assume that we have a (k,ϵ)-coreset S_{i-1} of P_{i-1}. Note that p_i is a (k,ϵ)-coreset of itself (in fact a $(k,0)$-coreset of itself), so $S_{i-1}\cup\{p_i\}$ is a (k,ϵ)-coreset of P_i. Let L be the cost of clustering $S_{i-1}\cup\{p_i\}$ to C. Therefore the cost of clustering P_i to C is in the interval $[(1-\epsilon)L,(1+\epsilon)L]$. Recall that the upper bound on $s(p)$ from Lemma 3.7 is:

$$\frac{\rho\bar{\alpha}D(p,p')}{\sum_{q\in P} D(q,q')} + \frac{\rho^2(\bar{\alpha}+1)}{|P(p)|}$$

By using L in place of the true cost $\sum_{p\in P_i} D(p,p')$ for defining $s'(p)$, the value of $t' = \sum_{p\in P_i} s'(p)$ increases to at most $\rho\bar{\alpha}\left(\frac{1+\epsilon}{1-\epsilon}\right)+\rho^2(\bar{\alpha}+1)\lambda k = O(\rho^4 k)$. Here there is no dependence on ϵ since we assume $\epsilon \leq 1/2$, so $\frac{1+\epsilon}{1-\epsilon}$ is bounded by an absolute constant.

2: Computing $|P(p)|$ is straightforward. Define $w(b) = |\{p\in P : \pi(p) = b\}|$ and then let $h : B \to C$ be the (γ,λ)-approximate clustering. Then $|P(p)| = \sum_{b\in h^{-1}(p')} w(b)$.

3: In Algorithm 1, we sample point p with probability $s'(p)$ to maintain a set M of non-deterministic size. We now argue that this can be converted to the desired coreset, where an i.i.d. sample of size m is taken from the distribution s'/t'. First, by a Chernoff bound, $|R| \geq E[|R|]/2$ with probability $1 - exp(-E[|R|]/8)$. Since $E[|R|] = xt' = (2c\epsilon^{-2}\log(\rho k)\log n\log(1/\delta))\cdot(\rho\bar{\alpha}+\rho^2(\bar{\alpha}+1)k) = \Omega(\log(n))$, we have that $|R| \geq E[|R|]/2$ with probability $1-O(1/n)$. Then by the union bound, this inequality holds true at each of the n iterations of receiving a point throughout the entire stream. Recall that for the offline coreset construction outlined in Subsection 3.3 to hold, we need an i.i.d. sample of at least $m = ct'\epsilon^{-2}(\log n\log t' + \log(1/\delta))$. By Lemma 3.7 $t' = \rho\bar{\alpha}+\rho^2(\bar{\alpha}+1)k$, and so by plugging in values we see that $E[|R|] = xt' \geq 2m$. Having that $|R| \geq m$ with probability $1-O(1/n)$,

it is well-known (see [DDH+07] for example) that this can be converted to the required i.i.d. sample.

□

4 Preliminaries for Offline Coreset Construction

4.1 Query space

In our framework, as in [FL11], we are given a finite input set P of items that are called *points*, a (usually infinite) set Q of items that are called *queries*, and a *cost* function f that maps each pair of a point in P and a query in Q to a non-negative number $f(p,q)$. The cost of the set P to this query is the sum over all the costs,

$$\bar{f}(P,q) := \sum_{p \in P} f(p,q).$$

More generally, each input point might be given a positive multiplicative weight $w(p) > 0$, and the overall cost of each point is then reweighed, so that

$$\bar{f}(P,w,q) = \sum_{p \in P} w(p) f(p,q).$$

The tuple (P,w,f,Q) is thus define our input problem and we call it a *query space*. In the case that the points are unweighted we can simply define $w(p) = 1$ for every $p \in P$.

However, for the following sections, it might help to scale the weights so that their sum is 1. In this case, we can think of the weights as a given distribution over the input points, and the cost $\bar{f}(P,w,q)$ is the expected value of $f(p,q)$ for a point that is sampled at random from P. For the unweighted case we thus have $w(p) = 1/n$ for each point, and

$$\bar{f}(P,w,q) = \frac{1}{n} \sum_{p \in P} f(p,q)$$

is the average cost per point. The cost function f is usually also scaled as will be explained later, to have values $f(p,q)$ between 0 and 1, or -1 to 1.

Example 1: Consider the following problem where the input is a set P of n points in $\mathbb{R}^d$. Given a ball $B = B(c,r)$ of radius r that is centered in $c \in \mathbb{R}^d$, we wish to compute the fraction of input points that are covered by this ball, $\frac{|P \cap B|}{|P|}$. More generally, the query is a set of k balls, and we wish to compute the fraction of points in P that are covered by the union of these balls. In this case, each input point $p \in P$ has a weight $w(p) = 1/n$, the set Q of queries is the union over every k balls in $\mathbb{R}^d$,

$$Q = \left\{ B(c_1,r_1) \cup \cdots \cup B(c_k,r_k) \mid \forall i \in [k] : c_i \in \mathbb{R}^d, r_i \geq 0 \right\}, \tag{2}$$

and the cost $f(p,q)$ for a query $q = \{B_1 \cup \cdots \cup B_k\} \in Q$ is either $1/n$ if p is inside one of the balls of q, and 0 otherwise. The overall cost $\bar{f}(P,w,q) = \sum_{p \in P} f(p,q)$ is thus the fraction of points of P that are covered by the union of these k balls.

The motivation of defining query spaces is usually to solve some related optimization problem, such as the query q that minimizes the cost $\bar{f}(P, w, q)$. In this case, the requirement to approximate every query in Q is too strong, and we may want to approximate only the optimal query in some sense. To this end, we may wish to replace the set Q by a function that assigns a different set of queries $Q(S)$ for each subset S of P. For the correctness of our results, we require that this function Q will be monotonic in the following sense: if T is a subset of S then $Q(T)$ must be a subset of $Q(S)$. If we wish to have a single set $Q(P)$ of queries as above, we can simply define $Q(S) := Q(P)$ for every subset S of P, so that the desired monotonic property will hold.

Example 2: For the set Q in (2), define $Q(S)$ to be the set of balls in $\mathbb{R}^d$ such that the center of each ball is a point in S. More generally, we can require that the center of each ball will be spanned by (i.e., linear combination of) at most 10 points from S.

We now conclude with the formal definitions for the above discussion.

Definition 4.1 (weighted set)**.** *Let S be a subset of some set P and $w : P \to [0, \infty)$ be a function. The pair (S, w) is called a* weighted set*.*

Definition 4.2 (query space)**.** [FL11]*. Let Q be a function that maps every set $S \subseteq P$ to a corresponding set $Q(S)$, such that $Q(T) \subseteq Q(S)$ for every $T \subseteq S$. Let $f : P \times Q(P) \to \mathbb{R}$ be a* cost function*. The tuple (P, w, Q, f) is called a* query space*.*

We denote the cost *of a query $q \in Q(P)$ by*

$$\overline{f}(P, w, q) := \sum_{p \in P} w(p) f(p, q).$$

4.2 (ε, ν)-Approximation

Consider the query space (P, w, Q, f) in Example 1 for $k = 1$, and suppose that we wish to compute a set $S \subseteq P$, such that, for every given ball B, the fraction of points in P that are covered by B, are approximately the same as the fraction of the points that are covered in S, up to a given small additive percentage error $\varepsilon > 0$. That is, for every ball B,

$$\left| \frac{|P \cap B|}{|P|} - \frac{|S \cap B|}{|S|} \right| \leq \varepsilon.$$

By defining the weight $u(p) = 1/|S|$ for every $p \in S$, this implies that for every query (ball) $q = B$,

$$\left| \bar{f}(P, w, q) - \bar{f}(S, u, q) \right| = \left| \sum_{p \in P} \frac{1}{|P|} \cdot f(p, q) - \sum_{p \in S} \frac{1}{|S|} \cdot f(p, q) \right| = \left| \frac{|P \cap B|}{|P|} - \frac{|S \cap B|}{|S|} \right| \leq \varepsilon.$$

A weighted set (S, u) that satisfies the last inequality for every query $q \in Q(S)$ is called an *ε-approximation* for the query space (P, w, f, Q). Note that the above example assumes that the maximum answer to a query q is $f(p, q) \leq 1$. Otherwise, the error guaranteed by an ε-approximation for a query q is $\varepsilon \max_{p \in P} |f(p, q)|$.

The above inequalities implies that if a ball covers a fraction of at least ε points from P (i.e., at least εn points), then it must cover at least one point of P. If we only ask for this (weaker) property from S then S is called an *ε-net*. To obtain the new results of this paper, we use a tool that generalizes the notion of ε-approximation and ε-net, but less common in the literature, and is known as *(ε,ν)-approximation* [LLS01].

By letting $a = \bar{f}(P,w,q)$ and $b = \bar{f}(S,u,q)$, an ε-approximation implies that $|a-b| \leq \varepsilon$ for every query q, and ε-net implies that $b > 0$ if $a \geq \varepsilon$. Following [LLS00], we define below a distance function that maps a positive real $|a-b|_\nu$ for each pair of positive real numbers a and b. A specific value ν will imply $|a-b|_\nu \leq \varepsilon$, i.e., that S is an ε-approximation, and a different value of ν will imply that S is an ε-net for P. This is formalized as follows.

Definition 4.3 ((ε,ν)-approximation [LLS00])**.** *Let $\nu > 0$. For every $a, b \geq 0$, we define the distance function*

$$|a-b|_\nu = \frac{|a-b|}{a+b+\nu}.$$

Let (P,w,Q,f) be a query space such that $\sum_{p\in P} w(p) = 1$ and $f : P \to [0,1]$. For $\varepsilon > 0$, a weighted set (S,u) is an (ε,ν)-approximation *for this query space, if for every $q \in Q(S)$ we have*

$$\left|\overline{f}(P,w,q) - \overline{f}(S,u,q)\right|_\nu \leq \varepsilon \tag{3}$$

Corollary 4.4 ([LLS00, HP11])**.** *Let $a, b \geq 0$ and $\tau, \nu > 0$ such that $|a-b|_\nu \leq \tau$. Let $\varepsilon > 0$. Then*

(i) (ε-sample/approximation). *If $\tau = \varepsilon/4$ and $\nu = 1/4$ then*

$$|a-b| \leq \varepsilon.$$

(ii) (ε-net). *If $\tau = 1/4$ and $\nu = \varepsilon$ then*

$$(a \geq \varepsilon \quad \Rightarrow \quad b > 0).$$

(iii) (Relative ε-approximation). *Put $\mu > 0$. If $\nu = \mu/2$ and $\tau = \varepsilon/9$ then*

(1) $a \geq \mu \quad \Rightarrow \quad (1-\varepsilon)a \leq b \leq (1+\varepsilon)a.$
(2) $a < \mu \quad \Rightarrow \quad b \leq (1+\varepsilon)\mu.$

4.3 Constructing ε-Approximations

Unlike the notion of coresets in the next section, the idea of ε-approximations is known for decades [VC71, Mat89]. In particular, unlike coresets, ε-approximations and (ε,ν)-approximations in general, can be constructed using simple uniform random sampling of the points. The size of the sample depends linearly on the complexity of the queries in the sense soon to be defined in this section.

Intuitively, and in most practical cases, including the examples in this paper, this complexity is roughly the number of parameters that are needed to define a single query. For

example, a ball in $\mathbb{R}^d$ can be defined by $d+1$ parameters: its center $c \in \mathbb{R}^d$, which is defined using d numbers and the radius $r > 0$ which is an additional number. A query of k balls is similarly defined by $k(d+1)$ parameters. However, by the set theory we have $|R| = |R^m|$ for every integer $c \geq 1$, which means that we can encode every m integers to a single integer. We can thus always reduce the number of parameters that are needed to define a query from m to 1 by redefining our cost function f and re-encoding the set of queries.

There are also natural examples of query spaces whose cost function f is defined by one parameter, but the size of the sampling needed for obtaining an ε-approximation is unbounded, e.g., the query space where $f(p,q) = \text{sign}(\sin(pq))$ and $P = Q = \mathbb{R}^d$, where $\text{sign}(x) = 1$ if $x > 0$ and 0 otherwise; see details e.g. in [**?**]. Hence, a more involved definition of complexity is needed as follows.

While the number of subsets from a given set of n points is 2^n, its can be easily verified that the number of subsets that can be covered by a ball in $\mathbb{R}^d$ is roughly $n^{O(d)}$. The exponent $O(d)$ is called the *VC-dimension* of the family of balls in $\mathbb{R}^d$. The following definition is a simple generalization by [FL11] for query spaces where the query set is a function and not a single set. The original definition of pseudo-dimension can be found e.g. in [LLS00] and is very similar to the definition of VC-dimension given by [VC71] as well as many other similar measures for the complexity of a family of shapes.

Definition 4.5 (dimension [FL11])**.** *For a query space (P, w, Q, f) and $r \in [0, \infty)$ we define*

$$\text{range}(q, r) = \{p \in P \mid w(p) \cdot f(p, q) \leq r\}.$$

The dimension *of (P, w, Q, f) is the smallest integer d such that for every $S \subseteq P$ we have*

$$\left| \{\text{range}(q, r) \mid q \in Q(S), r \in [0, \infty)\} \right| \leq |S|^d.$$

The main motivation of the above definition, is that it tells us how many samples we need to take uniformly at random from the input set P, to get an (ε, ν)-approximation (S, u) as follows.

Theorem 4.6 ([VC71, LLS00, FL11])**.** *Let (P, w, Q, f) be a query space of dimension d such that $\sum_{p \in P} w(p) = 1$ and $f : P \to [0, 1]$, and let $\varepsilon, \delta, \nu > 0$. Let S be a random sample from P, where every point $p \in P$ is sampled independently with probability $w(p)$. Assign a weight $u(p) = \frac{1}{|S|}$ for every $p \in S$. If*

$$|S| \in \Omega(1) \cdot \frac{1}{\varepsilon^2 \nu} \left(d \log\left(\frac{1}{\nu}\right) + \log\left(\frac{1}{\delta}\right) \right)$$

then, with probability at least $1 - \delta$, (S, u) is an (ε, ν)-approximation *for the query space (P, w, Q, f).*

5 Improved Coreset Framework

5.1 Improved (ε, ν)-approximations

In this section we show the first technical result of this paper: that the additive error $\varepsilon \max_{p\in P} |f(p,q)|$ in (3) can be replaced by ε, and the assumption $\sum_{i=1} w_i = 1$ in Theorem 4.6 can be removed. The condition for this to work is that the total importance value t, as defined below, is small. More precisely, the required sample size will be near-linear with t. Also, the uniform sampling in Theorem 4.6 will be replaced by non-uniform sampling with a distribution that is proportional to the importance of each point. This result is essentially a generalization and significant simplification of the framework in [FL11] that used ε-approximations for constructing coresets. Maybe more importantly: using the idea of (ε, ν)-approximations we are able to show that the sample size is near linear in t while in [FL11] it is quadratic in t. For some applications, such an improvement means turning a theoretical result into a practical result, especially when t is close to $\sqrt{n}$.

We define the *importance* of a point as the maximum absolute weighted cost $w(p)f(p,q)$ of a point p, over all the possible queries q, i.e,

$$s(p) := w(p) \max_{q\in Q(P)} |f(p,q)|,$$

and hope that this sum is small (say, constant or $\log n$), in other words, that not *all* the points are very important. More precisely, if the sum $t = \sum_{p\in P} s(p)$ of these costs is 1, then we prove below that a new query space (P, w', Q, f') can be constructed, such that an (ε, ν)-approximation (S, u) for the new query space would imply

$$\left||\overline{f}(P,w,q)| - |\overline{f}(S,u,q)|\right|_\nu \leq \varepsilon. \tag{4}$$

That is, the additive error in (3) is reduced as desired.

The new query space (P, w', Q, f') is essentially a re-scaling of the original weights by their importance, $w'(p) = s(p)$. To make sure that the cost of each query will still be the same, we need to define f' such that $w(p)f(p,q) = w'(p)f'(p,q)$. This implies $f'(p,q) := w(p)f(p,q)/s(p)$. While the new cost $\bar{f}'(P,w',q)$ is the same as the old one $\bar{f}(P,w,q)$ for every query q, the maximum value of $|f'(p,q)|$ is 1, by definition of $s(p)$, even if $|f(p,q)|$ is arbitrarily large. Hence, the additive error $\varepsilon|f'(p,q)|$ in (3) reduced to ε in (4).

More generally, an $(\varepsilon, \nu/t)$-approximation for (P, w', Q, f') would yield (4). Using the uniform sample construction of Theorem 4.6, this implies that to get (4) we need to increase the sample size by a factor that is nearly linear in t.

Theorem 5.1.

- *Let (P, w, Q, f) be a query space where $f(p,q) \leq 1$ for every $p \in P$ and $q \in Q(P)$.*
- *Let $s : P \to (0, \infty)$ such that $s(p) \geq w(p) \max_{q\in Q(P)} f(p,q)$.*
- *Let $t = \sum_{p\in P} s(p)$.*

- *Let* $w' : P \to [0,1]$ *such that* $w'(p) := s(p)/t$.
- *Let* $f' : P \times Q(P) \to [0,1]$ *be defined as* $f'(p,q) := w(p) \cdot \dfrac{f(p,q)}{s(p)}$.
- *Let* (S,u) *be an* $(\varepsilon, \nu/t)$-approximation *for* (P, w', Q, f').
- *Let* $u'(p) = u(p) \cdot \frac{w(p)}{w'(p)}$ *for every* $p \in S$.

Then for every $q \in Q(S)$,

$$|\overline{f}(P,w,q) - \overline{f}(S,u',q)|_\nu \leq \varepsilon. \tag{5}$$

Proof. Put $q \in Q(S)$.

$$\begin{aligned}
|\overline{f}(P,w,q) - \overline{f}(S,u',q)|_\nu &= \left| \sum_{p\in P} w(p)f(p,q) - \sum_{p\in S} \frac{u(p)w(p)}{\frac{s(p)}{t}} \cdot f(p,q) \right|_\nu && (6)\\
&= \left| t\sum_{p\in P} \frac{s(p)}{t} \cdot w(p) \cdot \frac{f(p,q)}{s(p)} - t\sum_{p\in S} u(p)\cdot w(p) \cdot \frac{f(p,q)}{s(p)} \right|_\nu \\
&= \left| t\sum_{p\in P} w'(p)\cdot f'(p,q) - t\sum_{p\in S} u(p)\cdot f'(p,q) \right|_\nu && (7)\\
&= \left| \sum_{p\in P} w'(p)\cdot f'(p,q) - \sum_{p\in S} u(p)\cdot f'(p,q) \right|_{\nu/t} && (8)\\
&\leq \varepsilon, && (9)
\end{aligned}$$

where (6) is by the definition of u', (7) is by the definition of f', (8) follows since

$$|ta - tb|_\nu = \frac{|ta-tb|}{ta+tb+\nu} = \frac{|a-b|}{a+b+\nu/t} = |a-b|_{\nu/t}, \tag{10}$$

and (9) is by (3) and the definition of S in the sixth bullet of the statement. □

Plugging Theorem 4.6 in Theorem 5.1 yields our main technical result that would imply smaller coresets in the next sections.

Theorem 5.2. *Let* (P,w,Q,f) *be a query space of dimension* d*, where* f *is non-negative, and let* $\varepsilon, \delta, \nu > 0$*. Let* S *be a random sample from* P*, where every point* $p \in P$ *is sampled independently with probability* $s(p)/t$*. Assign a weight* $u'(p) = \frac{tw(p)}{s(p)\cdot|S|}$ *for every* $p \in S$*. If*

$$|S| \in \Omega(1) \cdot \frac{t}{\varepsilon^2\nu}\left(d\log\left(\frac{t}{\nu}\right) + \log\left(\frac{1}{\delta}\right)\right)$$

then, with probability at least $1-\delta$,

$$|\overline{f}(P,w,q) - \overline{f}(S,u',q)|_\nu \leq \varepsilon.$$

Proof. By Theorem 5.1 it suffices to prove that S is an $(\varepsilon, \nu/t)$-approximation for (P, w', Q, f'). Indeed, since $\sum_{p\in P} w'(p) = 1$ and S is a random sample where $p \in P$ is sampled with probability $w'(p) = s(p)/t$, the weighted set (S, u) is, with probability at least $1 - \delta$, an $(\varepsilon, \nu/t)$-approximation for (P, w', Q, f') by Theorem 4.6. □

While Theorem 4.6 suggests to construct an ε-approximation simply by taking a uniform random sample from P, Theorem 5.2 requires us to take *non-uniform* sample where the distribution is defined by the importance $s(\cdot)$. Bounding these importances such that the sum $t = \sum_{p\in P} s(p)$ of importances will be small raise a new optimization problem that, as we will see in the next sections, might be not easy at all to solve. So what did we gain by proving Theorem 4.6 (or the framework of [FL11] in general)?

Most of the existing papers for constructing coresets essentially had to bound the total importance of the related problem anyway. However, without Theorem 4.6, their proofs also had to deal with complicated terms that involved ε and include sophisticated probability arguments. Essentially each paper had to re-prove in some sense the very involved proofs in [VC71, LLS00] that were researched for decades, as well as the mathematics behind the usage of Definition 4.3. Beside the complicated proofs, the final bounds, the dependency on ε and δ, as well as the coreset construction algorithm, were usually sub-optimal compared to Theorem 4.6. On the contrary, bounding the total importance allows us to focus on a deterministic results (no δ involved) and the terms $s(p)$ can be approximated up to constant factors (and not $(1 + \varepsilon)$-factors). This is demonstrated in Section 6.

5.2 Improved Coreset Framework

While the result in the previous section can be used to improve the quality of (ε, ν)-approximation in general, its main motivation is to construct the more recent type of data structures that is sometimes called coresets.

As explained in Theorem 4.6, ε-approximation and (ε, ν)-approximation in general, can be easily computed using uniform random sampling. While ε-approximations are useful for hitting sets, where we wish to know how much points are covered by a set of shapes, they are less relevant for shape fitting, where we wish to approximate the sum of *distances* of the input points to a given query shape or model. The main reason is that in shape fitting the maximum contribution of a point to the overall cost (sum of covered points) is bounded by 1, while in general its distance to the shape is unbounded. Using the notation of the previous section: the importance of each point is the same, so Theorem 5.2 yields uniform sample S and the same error as Theorem 4.6.

Example: In the Euclidean k-means problem the input is a set P of n points in $\mathbb{R}^d$. For simplicity, assume that the set is unweighted, that is $w(p) = 1/n$ for every $p \in P$. A query in this context is a set of points (centers) in $\mathbb{R}^d$ so $Q(P)$ is the family (set) of k balls in $\mathbb{R}^d$. We denote the squared Euclidean distance from a point $p \in P$ to a center $c \in \mathbb{R}^d$ by $D(p, c) = \|p - c\|_2^2$. The distance from p to its closest center in $C = \{c_1, \cdots, c_k\}$ is then

$$D(p, C) := \min_{c\in C} D(p, c) = \min_{c\in C} \|p - c\|_2^2 .$$

By defining $g(p, C) = D(p, C)$ we obtain the query space (P, w, g, Q), where $\bar{g}(P, w, C)$ is the average squared distances to a given (query) set C of k centers. Our goal is to compute a weighted subset (S, u) such that $\bar{g}(P, w, C)$ will approximate the average squared distances $\bar{g}(S, u, C)$ for every set of centers.

Suppose that (S, u) is an ε-approximation of (P, w, Q, g). Then

$$|\bar{g}(P, w, C) - \bar{g}(S, u, C)| \leq \varepsilon \max_{p \in P} g(p, C). \tag{11}$$

That is,

$$\left| \frac{1}{n} \sum_{p \in P} D(p, C) - \sum_{p \in P} u(p) D(p, C) \right| \leq \varepsilon \max_{p \in P} D(p, C). \tag{12}$$

In other words, the additive error depends on the maximum distance between a point to a given center, which can be arbitrary large, unless we assume, say, that both the input points and centers are inside the unit cube. Theorem 5.2 would not improve this bound by itself, since the importance of each point, $\max_{c \in Q(P)} D(p, C)$ is unbounded.

In this section we wish to compute a weighted subset (S, u) that will approximate the average distance $\bar{g}(P, w, C)$ for every query, up to a *multiplicative* factor of $1 \pm \varepsilon$ without further assumptions. Such a set is sometimes called a coreset as follows.

Definition 5.3 (ε-coreset)**.** *For an error parameter $\varepsilon \in (0, 1)$, the weighted set (S, u) is an ε-coreset for the query space (P, w, Q, g) if $S \subseteq P$ and for every $q \in Q(S)$,*

$$(1 - \varepsilon)\bar{g}(P, w, q) \leq \bar{g}(S, u, q) \leq (1 + \varepsilon)\bar{g}(P, w, q).$$

An equivalent definition of an ε-coreset, is that the additive error $\varepsilon \max_{p \in P} g(p, C)$ in (11) is replaced by $\varepsilon \bar{g}(P, w, q)$. That is, for every $q \in Q(S)$

$$|\bar{g}(P, w, q) - \bar{g}(S, u, q)| \leq \varepsilon \bar{g}(P, w, q).$$

ε-coreset implies that not only the average, but also the sum of squared distances (or sum of costs, in general) are preserved up to $1 \pm \varepsilon$. Note also that simply multiplying the weight of each point in S by $(1 - \varepsilon)$ would yield a one sided error,

$$\bar{g}(P, w, q) \leq \bar{g}(S, u, q)| \leq \frac{1 + \varepsilon}{1 - \varepsilon} \cdot \bar{g}(P, w, q) = 1 + \frac{2\varepsilon}{1 - \varepsilon} \cdot \bar{g}(P, w).$$

If we assume in addition that $\varepsilon \in (0, 1 - 2/c)$ for some $c > 2$, then $1 - \varepsilon \geq 2/c$ and thus

$$\bar{g}(P, w, q) \leq \bar{g}(S, u, q)| \leq (1 + c\varepsilon)\bar{g}(P, w).$$

Hence, an ε/c-coreset (S, u) implies,

$$\bar{g}(P, w, q) \leq \bar{g}(S, u, q) \leq (1 + \varepsilon)\bar{g}(P, w, q). \tag{13}$$

For example, if $\varepsilon \in (0, 1/2)$, then an $(\varepsilon/4)$-coreset would yield (13), by substituting $c = 4$.

The main observation for getting the desired multiplicative approximation is that a multiplicative approximation of $\varepsilon\bar{g}(P,w,q)$, can be turned into an additive approximation of ε by replacing $g(p,q)$ with its scaled version

$$f(p,q)=\frac{g(p,q)}{\bar{g}(P,w,q)}.$$

To get a coreset for g as in (13) it suffices to have an ε-approximation for f. In addition, for many problems, while the importance $s(p)$ of a point in p is unbounded with respect to g, it is bounded with respect to f. We formalize this observation as follows.

Corollary 5.4. *Let (P,w,Q,g) be a query space for some non-negative function g, and define $f:P\times Q(P)\to\mathbb{R}$ such that for every $p\in P$ and $q\in Q(P)$ we have*

$$f(p,q)=\frac{g(p,q)}{\bar{g}(P,w,q)}.$$

Let (S,u') be a $\left(\varepsilon/4,\frac{1}{4t}\right)$-approximation for (P,w',Q,f') as defined in Theorem 5.1. Then (S,u') is an ε-coreset for (P,w,Q,g), i.e., for every $q\in Q(S)$

$$|\bar{g}(P,w,q)-\bar{g}(S,u,q)|\leq\varepsilon\bar{g}(P,w,q).$$

Proof. Put $q\in Q(S)$, $\tau=\varepsilon/4$ and $\nu=1/4$. Applying Theorem 5.1 with ν yields

$$|\bar{f}(P,w,q)-\bar{f}(S,u',q)|_\nu\leq\tau.$$

Since $\bar{f}(P,w,q)=1$, this implies

$$\left|1-\frac{\bar{g}(S,u',q)}{\bar{g}(P,w,q)}\right|_\nu=|\bar{f}(P,w,q)-\bar{f}(S,u',q)|_\nu\leq\tau.$$

Substituting $a=1$, and $b=\bar{f}(S,u',q)$ in Corollary 4.4(i) yields

$$\left|1-\frac{\bar{g}(S,u',q)}{\bar{g}(P,w,q)}\right|\leq\varepsilon.$$

Multiplying by $\bar{g}(P,w,q)$ yields an ε-coreset as

$$|\bar{g}(P,w,q)-\bar{g}(S,u',q)|\leq\varepsilon\bar{g}(P,w,q).$$

□

Combining Corollary 5.4 and Theorem 4.6 yields the following theorem.

Theorem 5.5. *Let (P, w, Q, g) be a query space, where g is a non-negative function. Let $s : P \to (0, \infty)$ such that*

$$s(p) \geq \max_{q \in Q(P)} \frac{w(p)g(p,q)}{\sum_{p \in P} w(p)g(p,q)},$$

and $t = \sum_{p \in P} s(p)$. Let d be the dimension of (P, w', Q, f') as defined in Theorem 5.4. Let $c \geq 1$ be a sufficiently large constant, and let S be a random sample of

$$|S| \geq \frac{ct}{\varepsilon^2}\left(d \log t + \log\left(\frac{1}{\delta}\right)\right),$$

points from P, such that for every $p \in P$ and $q \in S$ we have $p = q$ with probability $s(p)/t$. Let $u'(p) = \frac{t \cdot w(p)}{s(p)|S|}$ for every $p \in S$. Then, with probability at least $1 - \delta$, (S, u') is an ε-coreset for (P, w, Q, g), i.e.,

$$\forall Q \in Q(S) : (1 - \varepsilon)\overline{g}(P, w, q) \leq \overline{g}(S, u', q) \leq (1 + \varepsilon)\overline{g}(P, w, q).$$

6 Application: Smaller and Generalized Coreset for k-Means

Definition 6.1 (ρ-metric space)**.** *Let X be a set, and $D : X^2 \to [0, \infty)$ be a function. Let $\rho > 0$. The pair (X, D) is a* ρ-metric *space if for every $(x, y, z) \in X^3$ we have*

$$D(x, z) \leq \rho(D(x, y) + D(y, z)).$$

For $C \subseteq X$ we denote $D(x, C) := \min_{c \in C} D(x, c)$ assuming that such minimum exists. Moreover, for $\varepsilon > 0$, the pair (X, D) is a (ψ, ε)-metric *space if for every $(x, y, z) \in X^3$ we have*

$$|D(x, z) - D(y, z)| \leq \frac{D(x, y)}{\psi} + \varepsilon D(y, z)).$$

Note that for every $x, y \in X$, and a center $c_y \in C$ that is closest to y, we have

$$D(x, C) = \min_{c \in C} D(x, c) \leq D(x, c_y) \leq \rho(D(x, y) + D(y, c_y)) = \rho(D(x, y) + D(y, C)). \quad (14)$$

A simple way to decide whether (D, X) is indeed a ρ-metric space is to use the following bound, known as Log-Log Lipschitz, that is usually easier to compute. The following lemma is very similar to [FS12, Lemma 2.1], where in the proof of (i) the constant 4 that appeared there is replaced here by $1/\varepsilon$.

Lemma 6.2 (Lemma 2.1(ii) in [FS12])**.** *Let $\tilde{D} : [0, \infty) \to [0, \infty)$ be a monotonic non-decreasing function that satisfies the following (Log-Log Lipschitz) condition: there is $r > 0$ such that for every $x > 0$ and $\Delta > 1$ we have*

$$\tilde{D}(\Delta x) \leq \Delta^r \tilde{D}(x).$$

Let (X, dist) *be a metric space, and let* $D(x, y) = \tilde{D}(\text{dist}(x, y))$ *for every* $x, y \in X$*. Then* (X, D) *is a* ρ*-metric space for* $\rho = \max\{2^{r-1}, 1\}$*, and a* (ψ, ε)*-metric space for every* $\varepsilon \in (0, 1)$ *and* $\psi = (r/\varepsilon)^r$*.*

For example, consider a metric space (X, dist) and the function $\tilde{D}(x) = x^2$ that corresponds to the squared distance $D(p, q) = \tilde{D}(\text{dist}(p, q)) = (\text{dist}(p, q))^2$. Note that (X, D) is not a metric space since the triangle inequality does not hold. However, for every $x > 0$ and $\Delta > 1$

$$\tilde{D}(x\Delta) = (x\Delta)^2 = \Delta^2 x^2 = \Delta^2 \tilde{D}(x).$$

Hence, by Lemma 6.2 $r = 2$, $\rho = 2^{r-1} = 2$, (D, X) is a 2-metric space and a (ψ, ε)-metric space for $\psi = (2/\varepsilon)^2$.

To compute a coreset for a problem we need to decide what are the important points, or more formally, to use Theorem 5.5 we need to bound the importance $s(p)$ of each point $p \in P$. To do this, we usually need to solve another optimization problem that is usually related to computing the query with the minimal cost. For example, the bound of importance of a point in the k-mean problem, as will be defined later, is based on the optimal solution for the k-means problem. Unfortunately, this optimization problem is usually hard and the main motivation for constructing the coreset in the first place.

There are two leeways from this chicken-and-egg problem:

(i) Use the merge-and-reduce approach that reduces the problem of computing a coreset for a large set of n items to the problem of computing coresets for $\frac{n}{2|S|}$ small weighted (core)sets. Each input coreset $2|S|$ is reduced to a coreset of $|S|$ and merged with another such coreset, where $2|S|$ is the minimum size of input set that can be reduced to half using the given coreset construction. In this case, if this coreset construction takes time $f(|S|)$ than, since there are such $O(n)$ constructions, the overall running time will then be $O(n) \cdot f(|S|)$.

(ii) For problems such as k-means, it is NP-hard to compute the optimal solution, even for a small set of $n = O(k)$ points. Instead of computing an optimal solution, usually a constant factor approximation suffices for computing the importance of each point. Since for many problems such an approximation is also unknown or too slow to compute, (α, β)-approximation, (or bi-criteria, or bicriterion), can be used instead as explained below.

Suppose that the optimal solution for the k-means problem on a set is OPT_k. That is, there is a set C of k centers whose cost (sum of squared distances to the input points of P) is OPT_k, and there is no such set of smaller cost. Then a set $\tilde{C}$ is called an α-approximation if its cost is at most $\alpha \cdot OPT$. However, for many problems, even a rougher approximation would do: instead of using k centers to approximate OPT_k by a factor of α, we use βk centers, where each input point may be assigned for the nearest center. Note that the cost is still compared to OPT_k and not to $OPT_{\beta k}$. We define (α, β)-approximation formally below. For our purposes later, we generalize this common definition of (α, β)-approximation, and

allow a point to be assigned to a different center than its nearest one, as long as the overall cost is small.

Definition 6.3 ((α,β)-approximation.)**.** *Let (X,D) be a ρ-metric space, and $k \geq 1$ be an integer. Let (P,w,Q,g) be a query space such that $P \subseteq X$, $Q(P) = \{C \subseteq X \mid |C| \leq k\}$, and $g : P \times Q(P)$ be a function such that*

$$g(p,C) = D(p,C) := \min_{c \in C} D(p,c). \tag{15}$$

Let $\alpha,\beta \geq 0$, $B \subseteq X$ such that $|B| \leq \beta k$ and $\mathcal{B} : P \to B$ such that

$$\sum_{p \in P} w(p) g(p, \{\mathcal{B}(p)\}) \leq \alpha \min_{q \in Q(P)} \bar{g}(P,w,q).$$

Then $\mathcal{B}$ is called an (α,β)-approximation *for (P,w,Q,g).*

For every $b \in B$, we denote by $P_b = \{p \in P \mid \mathcal{B}(p) = b\}$ the points that are mapped to the center b. We also denote $p' = \mathcal{B}(p)$ for every $p \in P$.

One of the main tools in our novel streaming algorithm is also a technique to update an (α,β)-approximation. However, due to memory limitations, our streaming algorithm cannot attach each point to its nearest center, but still the distances to the approximated centers is bounded. We thus generalize the definition of an (α,β)-approximation, which is usually a set B of size βk, to a *function* $\mathcal{B} : P \to B$ that assigns each input point to a (not necessarily closest) center in B, while the overall cost is still bounded.

Since an (α,β)-approximation yields a weaker result compared to a PTAS or α-approximation, it can usually be computed very quickly. Indeed, a very general framework for constructing (α,β)-approximation to any query space with a small VC-dimension is suggested in [FL11] where $\alpha = 1+\varepsilon$ and $\beta = O(\log n)$.

Reducing (α,β)-approximation to an $\alpha = O(1)$ approximation. The size of the coreset usually depends on α and β. However, if they are reasonably small (e.g. polynomial in $\log n$), we can reduce the approximation factor and number of centers in few phases as follows: (i) Compute an (α,β)-approximation for small (but maybe not constant) α and β. (ii) Compute an ε-coreset for $\varepsilon = 1/2$ using this approximation. (iii) Compute an $O(1)$ factor approximation on the coreset. Since the coreset is small, such an approximation algorithm can run inefficiently, say, in polynomial time if the coreset is of size $(\log n)^{O(1)}$. The resulting $O(1)$ approximation for the coreset is also an $O(1)$ approximation for the original set, by the definition of the $\varepsilon = 1/2$ coreset. (iv) Recompute the coreset for the complete (original) data using the $O(1)$ approximation instead of the (α,β)-approximation to obtain a coreset of size independent of both n and d.

Assumption 6.4. *In what follows we assume that:*

- *P is a set that is contained in X where (X,D) is a ρ-metric space and (ϕ,ε)-metric space as defined in* (6.1)*.*

Algorithm 2: CORESET$(P, w, \mathcal{B}, m)$

Input: A weighted set (P, w) where $P \subseteq X$ and (D, X) is a ρ-metric space, (α, β)-approximation $\mathcal{B} : P \to B$, and sample size $m \geq 1$.

Output: A pair (S, u) that satisfies Theorem 6.6.

for *each* $b \in B$ **do**
  Set $P_b \leftarrow \{p \in P \mid \mathcal{B}(p) = b\}$
for *each* $b \in B$ *and* $p \in P_b$ **do**
  Set $\text{Prob}(p) \leftarrow \frac{w(p)D(p,\mathcal{B}(p))}{2\sum_{q\in P} w(q)D(q,\mathcal{B}(q))} + \frac{w(p)}{2|B|\sum_{q\in P_b} w(q)}$.
Pick a sample S of at least m points from P such that for each $q \in S$ and $p \in P$ we have $q = p$ with probability $\text{Prob}(p)$.
for *each* $p \in P$ **do**
  Set $u(p) \leftarrow \frac{w(p)}{|S|\cdot\text{Prob}(p)}$.
Set $u(p) \leftarrow 0$ for each $p \in P \setminus S$. `/* Used only in the analysis. */`
return (S, u)

- (P, w, Q, g) *is a query space as defined in* (15).
- dk *denotes the dimension of* f' *as defined in Corollary 5.5.*
- c *is a sufficiently large constant that can be determined from the proofs of the theorems.*
- $\mathcal{B}$ *is an* (α, β)*-approximation for* (P, w, Q, g) *as in Definition 6.3.*
- *We are given an error parameter* $\varepsilon \in (0, 1)$.
- *We are given a maximum probability of failure* $\delta \in (0, 1)$.

We begin with the following claim, that is a simplified and generalized version of a similar claim in [LS10].

Lemma 6.5. *For every* $b \in B$*, and* $p \in P_b$ *have*

$$\frac{w(p)D(p,C)}{\sum_{q\in P} w(q)D(q,C)} \leq \frac{\rho\alpha w(p)D(p,p')}{\sum_{q\in P} w(q)D(q,q')} + \frac{\rho^2(\alpha+1)}{\sum_{q\in P_b} w(q)}.$$

Proof. Put $p \in P$ and $b \in B$ such that $p \in P_b$. We need to bound

$$\begin{aligned}\frac{w(p)D(p,C)}{\sum_{q\in P} w(q)D(q,C)} &\leq \frac{\rho w(p)D(p,p')}{\sum_{q\in P} w(q)D(q,C)} + \frac{\rho w(p)D(p',C)}{\sum_{q\in P} w(q)D(q,C)} \\ &\leq \frac{\alpha\rho w(p)D(p,p')}{\sum_{q\in P} w(q)D(q,q')} + \frac{\rho w(p)D(p',C)}{\sum_{q\in P} w(q)D(q,C)},\end{aligned} \tag{16}$$

where the first inequality holds by (14), and the second inequality holds since $\mathcal{B}$ is an (α, β)-approximation.

To bound the last term, we sum the inequality $D(p', C) \leq \rho(D(p', q) + D(q, C))$ over every $q \in P_b$ to obtain

$$\begin{aligned} D(p', C) \sum_{q \in P_b} w(q) &= \sum_{q \in P_b} w(q) D(p', C) \leq \sum_{q \in P_b} w(q) \cdot \rho(D(p', q) + D(q, C)) \\ &= \rho \sum_{q \in P_b} w(q) D(q', q) + \rho \sum_{q \in P_b} w(q) D(q, C) \\ &\leq \rho\alpha \sum_{q \in P_b} w(q) D(q, C) + \rho \sum_{q \in P_b} w(q) D(q, C) \\ &\leq \rho(\alpha + 1) \sum_{q \in P} w(q) D(q, C). \end{aligned}$$

Dividing by $\sum_{q \in P_b} w(q) \cdot \sum_{q \in P} w(q) D(q, C)$ yields

$$\frac{D(p', C)}{\sum_{q \in P} w(q) D(q, C)} \leq \frac{\rho(\alpha + 1)}{\sum_{q \in P_b} w(q)}.$$

Substituting this in (16) yields the desired result

$$\frac{w(p) D(p, C)}{\sum_{q \in P} w(q) D(q, C)} \leq \frac{\rho\alpha w(p) D(p, p')}{\sum_{q \in P} w(q) D(q, q')} + \frac{\rho^2(\alpha + 1)}{\sum_{q \in P_b} w(q)}.$$

□

Our first theorem suggests a coreset for k-means of size near-quadratic in k and quadratic in ε, based on our improved framework and last lemma. Existing work [LS10] for obtaining such coresets with only positive weights requires size cubic in k.

Theorem 6.6. *Under Assumption 6.4, let $t = k \cdot \rho^2(\alpha + 1)\beta$. Let (S, u) be the output of a call to algorithm* CORESET$(P, w, \mathcal{B}, m)$*, where*

$$m \geq \frac{ct}{\varepsilon^2} \left(dk \log t + \log \left(\frac{1}{\delta} \right) \right).$$

Then, with probability at least $1 - \delta$, (S, u) is an ε-coreset of size m for (P, w, Q, g).

Proof. Let $f : P \to [0, \infty)$ such that for every $C \in Q(P)$,

$$f(p, C) = \frac{D(p, C)}{\sum_{q \in P} w(q) D(q, C)},$$

and define

$$s(p) = \frac{\rho\alpha w(p) D(p, p')}{\sum_{q \in P} w(q) D(q, q')} + \frac{\rho^2(\alpha + 1)}{\sum_{q \in P_b} w(q)}.$$

By Lemma 6.5,

$$\sum_{p\in P}\max_{C\in Q(P)} w(p)|f(p,C)| = \sum_{b\in B}\sum_{p\in P_b}\max_{C\in Q(P)} w(p)f(p,C) \le \sum_{b\in B}\sum_{p\in P_b} s(p)$$
$$= \rho\alpha + |B|\cdot\rho^2(\alpha+1) = \rho(\alpha+|B|\alpha+|B|) \in O(\rho^2(\alpha+1)\beta k).$$

Applying Theorem 5.5 with the query space (P, w, Q, g) then yields the desired bound. □

Our second theorem in this section suggests a coreset for k-means of size near-linear in k by combining new observations with our improved framework.

Theorem 6.7. *Under Assumption 6.4, let $t = \alpha/\phi$. Let (S, u) be the output of a call to algorithm* CORESET$(P, w, \mathcal{B}, m)$*, where*

$$m \ge \frac{ck(t+\beta)}{\varepsilon^2}\left(d\log t + \log(\beta k) + \log\left(\frac{1}{\delta}\right)\right).$$

Then, with probability at least $1-\delta$, (S, u) is an ε-coreset of size m for (P, w, Q, g).

Proof. Let

$$H = \left\{p \in P \mid |D(p,C) - D(p',C)| \le \frac{2D(p,p')}{\phi}\right\}.$$

We need to bound by $O(\varepsilon)$ the expression

$$\frac{\left|\sum_{p\in P}(w(p)-u(p))\cdot D(p,C)\right|}{\sum_{q\in P} D(q,C)} \tag{17}$$

$$\le \frac{\left|\sum_{p\in H}(w(p)-u(p))\cdot(D(p,C)-D(p',C))\right|}{\sum_{q\in P} D(q,C)} \tag{18}$$

$$+ \quad \Big|\sum_{b\in B}\frac{\sum_{q\in P_b} w(q)D(b,C)}{\sum_{r\in P} w(r)D(r,C)}\cdot \tag{19}$$

$$\left(\sum_{p\in P_b}\frac{(w(p)-u(p))\cdot D(p,C)}{\sum_{q\in P_b} w(q)D(b,C)} - \sum_{p\in P_b\cap H}\frac{(w(p)-u(p))\cdot(D(p,C)-D(p',C))}{\sum_{q\in P_b} w(q)D(b,C)}\right)\Big|. \tag{20}$$

Put $b \in B$ and $p \in P_b$. Then,

$$\sum_{p\in P_b}\frac{(w(p)-u(p))\cdot D(p,C)}{\sum_{q\in P_b} w(q)D(b,C)} - \sum_{p\in P_b\cap H}\frac{(w(p)-u(p))\cdot(D(p,C)-D(p',C))}{\sum_{q\in P_b} w(q)D(b,C)}$$
$$= \sum_{p\in P_b}\frac{w(p)-u(p)}{\sum_{q\in P_b} w(q)} \tag{21}$$
$$+ \sum_{p\in P_b\setminus H}\frac{(w(p)-u(p))\cdot(D(p,C)-D(b,C))}{\sum_{q\in P_b} w(q)D(b,C)}. \tag{22}$$

We now prove that, with probability at least $1-c\delta$, each of the expressions (18), (21) and (22) is bounded by 2ε.

Bound on (18)**:** Let $h(p,C) = \frac{D(p,C)-D(p',C)}{\sum_{q\in P} w(q)D(q,C)}$ if $p \in H$ and $h(p,C) = 0$ otherwise. For every $p \in P$,

$$w(p)\cdot|h(p,C)| \leq \frac{2w(p)D(p,p')}{\phi\sum_{q\in P} w(q)D(q,C)} \leq \frac{2\alpha w(p)D(p,p')}{\phi\sum_{q\in P} w(q)D(q,q')} \leq \frac{2\alpha\text{Prob}(p)}{\phi}.$$

Hence, using $t = 2\alpha/\phi$ in Theorem 5.5, with probability at least $1-\delta$,

$$\frac{\left|\sum_{p\in H}(w(p)-u(p))\cdot(D(p,C)-D(p',C))\right|}{\sum_{q\in P} D(q,C)} = \left|\sum_{p\in H}(w(p)-u(p))h(p,C)\right| = |\bar{h}(P,w,q)-\bar{h}(S,u,q)| \leq \varepsilon, \tag{23}$$

Bound on (21)**:** Let $I(p,b) = 1/\sum_{q\in P_b} w(q)$ if $p \in P_b$ and $I(p,b) = 0$ otherwise. We have

$$\max_{p\in P} w(p)I(p,b) \leq 2|B|\text{Prob}(p),$$

where Prob(p) is defined in Line 3 of Algorithm 2. Hence, $\sum_{p\in P}\max w(p)I(p,b) \leq 2|B|$. Also, each point in S is sampled with probability proportional to $2|B|\text{Prob}(p)$ and for $c \geq 2$,

$$|S| \geq \frac{2|B|}{\varepsilon^2}\left(2\log(2|B|)+\log\left(\frac{1}{\delta}\right)\right) = \frac{2|B|}{\varepsilon^2}\left(\log(|B|)+\log\left(\frac{|B|}{\delta}\right)\right).$$

Using $d = 1$ and replacing δ by $\delta/|B|$ in Theorem 5.5 yields that S is an ε-coreset for $(P,w,\{b\},I)$, with probability at least $1-\delta/|B|$. That is

$$|\sum_{p\in P_b}\frac{w(p)-u(p)}{\sum_{q\in P_b} w(q)}| = |\bar{I}(P,w,b)-\bar{I}(S,u,b)| \leq \varepsilon. \tag{24}$$

By the union bound, with probability at least $1-\delta$, the last inequality holds for every $b \in B$ simultaneously.

Bound on (22)**:** Since (D,X) is a ρ-metric,

$$|D(p,C)-D(b,C)| \leq \frac{D(p,b)}{\phi}+\varepsilon D(b,C) \leq \max\left\{\frac{2D(p,b)}{\phi}, 2\varepsilon D(b,C)\right\}.$$

Hence

$$\sum_{p\in P_b\setminus H}\frac{(w(p)-u(p))\cdot(D(p,C)-D(b,C))}{\sum_{q\in P_b} w(q)D(b,C)} \leq 2\varepsilon\sum_{p\in P_b\setminus H}\frac{w(p)+u(p)}{\sum_{q\in P_b} w(q)}. \tag{25}$$

The last expression is bounded using (24), as

$$\sum_{p\in P_b\setminus H}\frac{w(p)+u(p)}{\sum_{q\in P_b} w(q)} \leq \sum_{p\in P_b\setminus H}\frac{(1+\varepsilon)w(p)}{\sum_{q\in P_b} w(q)} = (1+\varepsilon). \tag{26}$$

Bounding (17)**:** By combining the above bounds we obtain that with probability at least $1 - 10\delta$,

$$\frac{\left|\sum_{p\in P}(w(p) - u(p)) \cdot D(p, C)\right|}{\sum_{q\in P} D(q, C)} \leq \varepsilon + \left|(\varepsilon + 2\varepsilon(1+\varepsilon)) \sum_{b\in B} \frac{\sum_{q\in P_b} w(q)D(b, C)}{\sum_{r\in P} w(r)D(r, C)}\right| \leq c\varepsilon.$$

Replacing ε with ε/c, and δ with δ/c then proves the theorem.

□

The values of ρ and ψ. Algorithm 2 can be applied to compute a coreset for any given variant of the k-means/median problem given a set P and a ρ or (ρ, ε)-metric (X, D). The only difference is the size of the required coreset. The parameters ρ and ϕ can usually be computed easily using Lemma 6.2. For example, in the case of distances to the power of $r \geq 1$, the value of ρ is roughly 2^r and the value of ϕ is roughly ε^r. For most common m-estimators the values are similar, when r is some constant.

The values of α and β can be $\alpha = \beta = O(1)$ for k-means and all its variants, by using the generic algorithm for computing (α, β)-approximation in [FL11] with $\alpha = \beta = O(\log n)$, and then use the technique for reducing α and β. For bounding the approximation of the bi-criteria approximation in [FL11] only the pseudo-dimension of the problem is required, which is usually easy to compute as explained below.

Dimension d. Unlike the total sensitivity t, the dimension d for numerous problems was already computed in many papers in computational geometry and machine learning (in the context of PAC learning). This includes reductions and connections to similar notions such as the shattering dimension or the VC-dimension of a set. General techniques for computing the dimension of a set based on number of parameters to define a query, or number of operations that are needed to answer a query can be found in the book [AB99].

Dimension of the k-means problem and its variants. Note that, unlike sensitivity, the dimension is less dependent on the exact type of distance function. For example, using Euclidean distance or Euclidean distance to the power of 3 as a variant for the k-means clustering problem does not change the dimension. This is because set of ranges for both of these problems is the same: subsets of the input points that can be covered by k balls. It is easy to compute the dimension for the k-means problem (the query space (P, w, Q, g) in our paper, as well as the modified query space for the function f. These bounds can be found in [FL11]. In addition, [FL11] provide a simple reduction that shows that the dimension for k centers is the same as the dimension of 1 center multiplied by k.

In short, for the Euclidean space, the dimension of the k-means/median problem is $O(dk)$, and for metric spaces (graph) the dimension is $O(k \log n)$.

Smaller coreset for k-means queries. Consider the k-means queries in $\mathbb{R}^d$, i.e., the cost is the sum of squared distances $\sum_{p\in P} D(p, C)$ over every point in P to its nearest center in a given set C of k points in $\mathbb{R}^d$. It was proven that projecting P onto an $O(k/\varepsilon)$-dimensional subspace that minimizes its sum of squared distances, known as the low-rank approximation of P, would preserve this sum, to any set of k centers, up to a factor of $1 \pm \varepsilon$. This is in some sense an ε-coreset for P of size n and that is not subset of the input, but of low-dimensionality. In particular, this result implies that there is a set of centers (known as centroid set [HPK07]) that is contained in a $O(k/\varepsilon)$-dimensional space, such that every set of k centers in $\mathbb{R}^d$ can be replaced by a set of k centers in the centroid set, that would yield the same cost up to a factor of $1 \pm \varepsilon$. In particular, this implies that the dimension of the k-means problem can be reduced to $O(k/\varepsilon)$ instead of dk, i.e., independent of d. Combining this result with our paper yields the first coreset for k-means of size independent of d that is subset of the input (in particular, preserve the sparsity of the input points) that also supports streaming.

Weak coresets of size independent of d. For the non-Euclidean case or non-squared distances it seems that it is impossible to obtain coreset of size independent of d. However, coreset as defined above (sometimes called strong coreset) approximates *every* query in the set of queries, while the main application and motivation for constructing coreset is to compute the *optimal* query or its approximation. A weak coreset is a small set that can be used to give such an approximation. The exact definition of weak coreset also changes from paper to paper. In particular, a weak coreset for k-means was suggested in [FMS07]. However, to extract the approximated solution from the coreset we must run an exhaustive search and cannot use existing algorithms or heuristics as in the case of strong coreset.

In this paper, following [FL11], we use a simple and general definition of weak coreset, that is also more practical. Instead of defining a unique (static, global) set Q of queries, we define Q to be a function that maps every subset (potential coreset) S of P to a set of queries. A weak coreset S needs only to approximate the queries in $Q(S)$. It turns out that for many case the (generalized definition) of dimension for such a query space is much smaller compared to the traditional case where $Q(S) = Q(P)$ is the same for every subset $S \subseteq P$.

To be able to use this property in our existing proofs, we require a monotonicity property: that the set $Q(T)$ of queries that are assigned to a subset $T \subset S$ must be contained in $Q(S)$. If we can prove that for every $S \subseteq P$, the set $Q(S)$ contains a $(1+\varepsilon)$-approximation to both the optimal solution of S and P, then we can extract such an approximation from S. For k-means, k-median and their variants, it was proven in [SV07] that the optimal k-centers of a set S can be approximated by such a set that is spanned by $O(k/\varepsilon)$ points in S. By defining $Q(S)$ to be the union of the optimal center of P, with all the centers that are spanned by $O(k/\varepsilon)$ points in S, we get that S is a weak coreset. Note that the definition of Q can be explicit (without knowing the optimal center of P) and is needed only to bound its dimension as in Definition 4.5.

Extracting a $(1+\varepsilon)$-approximation from the coreset. It was proven in [FL11] that for problems such as k-median such weak coresets have dimension $O(k/\varepsilon)$, i.e., independent of d. Unlike [FL11], we suggest here a very simple way to extract the approximated solution from the coreset: compute the weighted coreset (S, u) (of size independent of d) as defined in Algorithm 2, and then use any given algorithm to compute a $(1+\varepsilon)$ approximation set C of k centers on the coreset (or any other trusted heuristic that we hope computes such a set C). Since C is not necessarily spanned by few points in S, it may not be a good approximation for the original set P and we should not return it. However,the proof in [SV07] is constructive and shows a near-linear time algorithm that using such an approximated solution set C, we can compute another $(1+O(\varepsilon))$-approximation C' that has the additional property that C' is spanned by $O(k/\varepsilon)$ points in S. Hence, C' is both a near-optimal solution to S and in $Q(S)$, so it must be a $(1+\varepsilon)$-approximation for the optimal solution of P.

7 Appendix A: Merge and Reduce Tree

We now briefly introduce the previous technique for maintaining coresets in the streaming setting due to Har-Peled and Mazumdar [HPM04] and Bentley and Sax [BS80]. In this method, a merge-and-reduce tree is built by using an offline coreset construction as a black-box. Previously merge-and-reduce was the only known technique for building a streaming coreset for metric k-median, and it relies solely on the following two properties:

1. Merge: The union of (k, ϵ)-coresets is a (k, ϵ)-coreset.
2. Reduce: A (k, ϵ)-coreset of a (k, δ)-coreset is a $(k, \epsilon + \delta)$-coreset.

The merge-and-reduce tree works as follows. There are buckets B_i for $i \geq 0$. In each step, the bucket B_0 takes in a segment of $O(1)$ points from the stream. Then the tree works like counting in binary: whenever buckets B_0 to B_{i-1} are full, these i buckets are merged and then reduced by taking a $(k, \frac{\epsilon}{\log n})$-coreset and storing the result in B_i.

Let s be the space of offline construction, which depends on ϵ as ϵ^{-a}. At the end of the stream, $O(\log n)$ buckets have been used and each bucket uses $O(s \log^a n)$ space; this incurs a multiplicative overhead of $\Theta(\log^{a+1} n)$ in the storage requirement. The second factor comes from using the accuracy parameter $\frac{\epsilon}{\log n}$, which is necessary by Property 2 since the construction will be compounded $O(\log n)$ times. Due to this compounding, the runtime is multiplied by a factor of $O(\log n)$.

8 Appendix B: General Streaming Reduction

We present a general technique for converting an offline coreset construction to a streaming coreset construction with $O(\log n)$ overhead. Given a ρ-metric space (X, D) (recall Definition 6.1), we build a query space (P, w, g, Q) in the same way as in Definition 6.3: $Q(P) = \{C \subseteq X \mid |C| \leq k\}$ and $g(p, C) = D(p, C) := \min_{c \in C} D(p, c)$. Here k is a positive integer that denotes how many centers may be used for the clustering.

Our bicriterion algorithm is an adjustment to the algorithms of [BMO+11] and [CCF] with the following important difference: our bicriterion is online, so we do not delete and reassign centers as is [BMO+11]. This "online" property is critical for the algorithm to work and is one of the main technical ideas. Although a fully online bicriterion can require linear space, we maintain a division of the stream P into a prefix R and a suffix $P \setminus R$ such that our bicriterion is online on the suffix $P \setminus R$ and the prefix R can be largely ignored. To maintain this property of being online for the suffix, we incur only a modest space increase from $O(k \log n)$ to $O(\log(\frac{1}{\epsilon})k \log n)$.

After having an online bicriterion (it is further explained below why this property is essential), the offline coreset algorithms perform non-uniform sampling procedure with carefully chosen probabilities that are defined by the bicriterion. Equipped with our new bicriterion algorithm, implementing the sampling procedure is rather straightforward computation which is explained in Section 8.2. As a result, we can implement any sampling-based coreset algorithm for k-median without merge-and-reduce and in one pass. As such it is applicable to several coreset constructions (such as k-means and other M-estimators). In addition, we believe that our methods will work with other objective functions as well such as (k, j)-subspace, and we hope that future work will investigate these directions.

Many clustering algorithms [COP03, GMM+03, BMO+11] maintain a weighted set (B, u) of points (which are selected using a facility-location algorithm). Upon arrival of an update $(p, w(p))$ from the stream, this update is added to the set by $u(p) \leftarrow u(p) + w(p)$.

In these algorithms, only a single operation is performed on (B, u) which we call MOVE. For two points $p, p' \in B$ with weights $u(p)$ and $u(p')$, the function $\mathsf{MOVE}(p, p')$ does the following: $u(p') \leftarrow u(p') + u(p)$ and $u(p) \leftarrow 0$. This essentially moves the weight at location p to the location of p'. The motivation for building (B, u) will be compression; (B, u) will be maintained over the stream P in such a way that $|B| = O(\log |P|)$.

Throughout this section, we will assume for ease of exposition that each point in the stream is from a distinct location. This simplifies the analysis, allowing $\mathcal{B}$ to take only a single value for each input point. The algorithm works for general inputs without modification, requiring only a more careful notation for the analysis. Additionally, we state the algorithm for unweighted input (where $w(p) = 1$ for all $p \in P$) and the parameter n is defined as $|P|$. We still include $w(p)$ throughout the algorithm and analysis, as the analysis generalizes to weighted inputs where the parameter n is replaced by the sum of all weights (after normalizing the minimally weighted point to 1).

8.1 An Algorithm for building a coreset

First, let us describe how the algorithm of [BMO+11] works. We will modify this algorithm as part of our coreset construction. In this summary, we alter the presentation from that of [BMO+11] to more fluidly transition to our modified version but the algorithm remains the same. [BMO+11] operates in phases $i \geq 1$. This means that the algorithm maintains a phase number i (used internally by the algorithm), beginning in phase $i = 1$. As the stream arrives, the algorithm may decide to increment the phase number. Let (R_i, w_i) denote the prefix of the input received before the end of phase i, and let $\mathsf{OPT}_k(R_i)$ denote the minimal value of

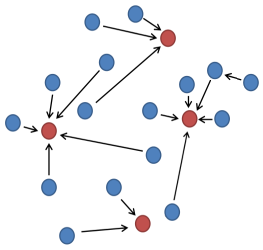

Figure 1: A bicriterion approximation. Although connections are not optimal, the sum of all connections costs is $O(\mathsf{OPT})$.

$\bar{g}(R_i, w_i, C)$ over every $C \in Q$. When phase $i+1$ begins, a value L_{i+1} is declared on Line 27 as a lower-bound for the current value of $\mathsf{OPT}_k(R_{i+1})$. The algorithm has computed (M_i, u_i), which we inductively assume is a bicriterion approximation for R_i (more precisely, a map $\mathcal{B} : R_i \to M_i$ such that $\cup_{p \in R_i}(\mathcal{B}(p), w(p)) = (M_i, u_i)$. However, to maintain polylogarithmic-space the algorithm pushes (M_i, u_i) to the beginning of the stream and restarts the bicriterion construction. This means that the algorithm, at this point, restarts by viewing the stream as $(P, w - w_i + u_i)$ (i.e. replacing (R_i, w_i) with (M_i, u_i)). Continuing in this way, the algorithm maintains a bicriterion (M_{i+1}, u_{i+1}) for $(R_{i+1}, w_{i+1} - w_i + u_i)$ (which is also a bicriterion for (R_{i+1}, w_{i+1}) by Theorem 8.2) until the next phase change is triggered.

Now we explain our modifications to [BMO+11] (see Algorithm 3). The first step is that the bicriterion our algorithm builds must be determined "online" in the following sense: upon receiving a point $(x, w(x))$, the value of $\mathcal{B}(x)$ must be determined (and never be altered) before receiving the next point from the stream.

This generalization is necessary for the following reason. Suppose we connect an element p to a center b_1. Later in the stream, we open a new center b_2 that becomes the closest center to p. However, using polylogarithmic space, we have already deleted b_1 and/or p from memory and the state of our algorithm is identical to the case where b_1 remains the closest center to p. Therefore the connections must be immutable, and this results in non-optimal connections.

Definition 8.1. *An online $[\alpha, \beta]$-bicriterion is an algorithm that maintains a bicriterion (B, t) over a stream X, and operates online in the following sense. Upon arrival of each point p, it is immediately decided whether p is added to B, and then $t(p)$ is determined. Both of these decisions are permanent.*

Upon receiving an update $(p, w(p))$ from the stream, the algorithm may call $\mathsf{MOVE}(p, p')$ for the nearest $p' \in B$ to p. In the analysis we use the function $\mathcal{B} : P \to B$ that maps each point p to its immediate location after this initial move (either p itself, or p'). If future moves are performed on p, this does not change the value $\mathcal{B}(p)$. $\mathcal{B}$ is not stored by the algorithm due to space constraints; only the value of $\mathcal{B}(p)$ for the most recent point p is used by the algorithm, and older values are used in the analysis only. We will show that $\mathcal{B}$ is a $(O(1), O(\log n))$-approximation that allows us to maintain a coreset over the stream.

We now state Algorithm 3. ϕ and γ are constants (dependent on ρ) used in the analysis that are defined as in [BMO+11]. Each point has a flag that is either raised or lowered (given by the *Flag* function). All points have their flag initially lowered, given on Line 1. A lowered flag shows that the point is being read for the first time (being received from the stream), and a raised flag shows that the point is being re-read by the algorithm (having been stored in memory).

On Line 22, (M_i, u_i) is the weighted set (B, u) as it exists at the end of phase i. We define the cost of $\mathsf{MOVE}(p, p')$ to be $w(p)D(p, p')$. The value K_i is therefore the total cost of all moves performed in phase i.

At a phase change, the set (B, u) is pushed onto the beginning of the stream with all points having a raised flag. This means that during the next $|B|$ iterations of the outer-loop where a point $(x, w(x))$ is received we actually receive a point from memory. We continue to process the stream after these $|B|$ points are read.

The following theorem summarizes the guarantees of this algorithm. We note that, as stated, the algorithm's runtime of $O(nk \log n)$ is not correct - in fact the number of phases may be arbitrarily large. However, using the same technique as detailed in Section 3.3 of [BMO+11] the number of phases can be bounded to $O(\frac{n}{k \log n})$ while requiring $O(k^2 \log^2 n)$ time per phase.

Theorem 8.2 ([BMO+11])**.** *Let Algorithm 3 process a stream (P, w) of at most n points. Let R_i denote the prefix of the stream received before the end of phase i, and let K_i and L_i be their final values from the algorithm (which are never modified after phase i). With probability at least $1 - \frac{1}{n}$, the following statements all hold after processing each point:*

1. $\sum_i K_i \leq \frac{\rho\phi\gamma}{\phi - \rho} \mathsf{OPT}_k(R_i)$
2. *The total runtime is $O(nk \log n)$*
3. *For every phase i, $L_i \leq \mathsf{OPT}_k(R_i) \leq \phi L_i \leq L_{i+1}$*
4. *At the execution of Line 22, M_i consists of $O(k \log n)$ points*

Part 1 of the preceding theorem, which bounds the total cost of all moves performed by the algorithm, also serves as an upper bound on the Earth-Mover distance between the stream (P, w) and the maintained set (B, u).

Definition 8.3 (Earth-Mover Distance)**.** *Let (A, w_A) and (B, w_B) be weighted sets in (X, D). Morever, let them be of equal weight in the sense that $\Sigma_{a \in A} w(a) = \Sigma_{b \in B} w(b)$. Define a "movement" from (A, w_A) to (B, w_B) to be a weighted set (Y, v) in $(X \times X, D)$ such that $\cup_{(a,b) \in Y}(a, v(a, b)) = (A, w_A)$ and $\cup_{(a,b) \in Y}(b, v(a, b)) = (B, w_B)$. Then $d_{EM}(A, w_A, B, w_B)$ is the minimum of $\Sigma_{(a,b) \in Y} v((a, b))D(a, b)$ over all movements Y from (A, w_A) to (B, w_B).*

Another way to view the preceding definition is in terms of probability distributions. Over all joint distributions over $A \times B$ with marginal distributions (A, w_A) and (B, w_B), we seek the minimum possible cost (as defined) for any such joint distribution.

Algorithm 3: Input: integer k, ρ-metric space (X, D), stream (P, w) of n weighted points from (X, D). Output: $\mathcal{B}(x)$ after receiving each point x, and a weighted set (M_i, u_i) after each phase $i \geq 1$

$Flag(x) = 0$ for all $x \in P$
$L_1 \leftarrow$ minimum $D(x, y)$ for any x, y in the first k distinct points
$i \leftarrow 1$
$K_1 \leftarrow 0$
$B \leftarrow \emptyset$
$u(x) \leftarrow 0$ for all x
for *each point* $(x, w(x))$ *received* **do**
    $u(x) \leftarrow u(x) + w(x)$
    $y \leftarrow argmin_{y \in Q} D(x, y)$ (break ties arbitrarily)
    $I \leftarrow 1$ with probability $\min\{\frac{w(x)D(x,y)}{L_i/k(1+\log_2 n)}, 1\}$, otherwise $I \leftarrow 0$
    if I **then**
        if $Flag(x) = 0$ **then**
            $\mathcal{B}(x) \leftarrow x$
            $B \leftarrow B \cup \{x\}$
    else
        $K_i \leftarrow K_i + w(x)D(x, y)$
        $u(y) \leftarrow u(y) + u(x)$ /* Step 1 of MOVE(x, y) */
        $u(x) \leftarrow 0$ /* Step 2 of MOVE(x, y) */
        if $Flag(x) = 0$ **then**
            $\mathcal{B}(x) \leftarrow y$
    if $K_i > \gamma L_i$ *or* $|B| > (\gamma - 1)(1 + \log_2 n)k$ **then**
        $(M_i, u_i) \leftarrow (B, u)$
        $Flag(b) \leftarrow 1$ for all $b \in B$
        Push (B, u) onto the stream (P, w) before the next point to read
        $B \leftarrow \emptyset$
        $u(x) \leftarrow 0$ for all x
        $L_{i+1} \leftarrow \phi L_i$
        $q \leftarrow 0$
        $K_{i+1} \leftarrow 0$
        $i \leftarrow i + 1$

From now on we will write a weighted set A instead of (A, w_A) when the meaning is clear. If we start with a set A_0 and apply n operations of MOVE until it becomes the set A_n, we can provide an upper-bound for $d_{EM}(A_0, A_n)$ by summing $d_{EM}(A_i, A_{i+1})$ for $0 \leq i < n$. This is a direct application of the triangle inequality. And if A_{i+1} is obtained from A_i by applying $\mathsf{MOVE}(p, p')$, then $d_{EM}(A_i, A_{i+1}) = D(p, p')w(p)$, the cost of this move.

The Earth-Mover distance is important for clustering problems for the following reason. For any weighted sets (P, w) and (B, u) and query C, $|\bar{g}(P, w, C) - \bar{g}(B, u, C)| \leq d_{EM}(P, w, B, u)$. This is immediate from a repeated application of the triangle-inequality (proofs are found in Theorem 2.3 of [GMM+03] as well as in [COP03, BMO+11, Guh09]).

Theorem 8.4. *There exists an algorithm stores $O(\log\left(\frac{1}{\epsilon}\right) k \log n)$ points and maintains for every prefix (R, w') of the stream (P, w): (1) a weighted set (M, u) such that $d_{EM}(M, u, R, w') \leq \epsilon \mathsf{OPT}_k(R)$, and (2) a $(O(1), O(\log(\frac{1}{\epsilon}) \log n))$-bicriterion $\mathcal{B}$ for (R, w'). Morever, this bicriterion $\mathcal{B}$ is computed online in the sense that $\mathcal{B}(x)$ is determined upon receiving x from the stream.*

Proof. The algorithm will consist of running Algorithm 3 and storing certain information for the λ most recent phases (where λ depends on ϵ).

Let i be the current phase number, and let (R, w') be the points received so far. We remind the reader that when the meaning is clear we suppress notation and write R for the weighted set (R, w'), and likewise for (M, u). Define $\lambda = 2 + \lceil \log_\phi(\frac{\rho\gamma}{(\phi-\rho)\epsilon}) \rceil$. The prefix R and the set M in the statement of the theorem will be $R_{i-\lambda}$ and $M_{i-\lambda}$. To upper bound $d_{EM}(R, M)$, which by definition is the minimum cost of any movement from R to M, we note that one movement is the set of moves carried out by the algorithm through phase $i - \lambda$, whose cost is $\sum_{j=1}^{i-\lambda} K_j$. Part 1 of Theorem 8.2 then shows that $d_{EM}(R, M) \leq \Sigma_{j=1}^{i-\lambda} K_j \leq \frac{\rho\phi\gamma}{\phi-\rho} \mathsf{OPT}_k(R_{i-\lambda})$. Combining Statements 3 and 4 of the theorem shows that $\mathsf{OPT}_k(R_{i-\lambda}) \leq \phi^{1-\lambda} \mathsf{OPT}_k(R)$. Therefore $d_{EM}(R, M) \leq \frac{\rho\phi\gamma}{\phi-\rho} \phi^{1-\lambda} \mathsf{OPT}_k(R) \leq \epsilon \mathsf{OPT}_k(R)$ as desired.

As for the second statement of the theorem, the algorithm defines the map $\mathcal{B}$ and we are currently interested in the restriction of $\mathcal{B}$ to $R \setminus R_{i-\lambda}$. $\mathcal{B}$ maps to at most $O(k \log n)$ points per phase - this is guaranteed by Statement 5 (a direct result the termination condition on Line 17). Over the last λ phases, this then maps to $O(\lambda k \log n) = O((\frac{1}{\epsilon}) k \log n)$ points. Therefore we have that $\beta = O((\frac{1}{\epsilon}) \log n)$. And the value of α is immediate from Statement 1 of Theorem 8.2, since $\mathcal{B}$ incurs only a subset of the costs as K_j (the subset that comes from the new portion of the stream $R_j \setminus R_{j-1}$).

As a final note, we do not store $\mathcal{B}(P)$ in memory (it may be linear in space). For algorithms in the following sections it is only required to know it's value for the most recent point received; previous values are used only in the analysis. □

8.2 Maintaining a coreset over the stream

In the previous section, we presented an algorithm that maintains an (α, β)-approximation of the stream (this is given by the function $\mathcal{B} : P \to B$). In this section we show how we can use this approximation to carry out the coreset construction of Algorithm 2 on an insertion-only

stream. In the offline construction, a sample S of m points is taken from X according to a distribution where point p sampled with probability depending on $D(p, \mathcal{B}(p))$, $|B|$, and $n_{\mathcal{B}(p)}$ (the total weight of points connected to the center $\mathcal{B}(p)$, which is written as $\Sigma_{q \in P_b} w(q)$ where $b = \mathcal{B}(p)$) - the specific formula for the probability written below (and comes from Line 3 of Algorithm 2). All three of these quantities can easily be maintained over the stream using Algorithm 4 since $\mathcal{B}$ is online (i.e. $\mathcal{B}(p)$ never changes).

$$Prob(p) = \frac{w(p)D(p, \mathcal{B}(p))}{2\sum_{q \in P} w(q)D(q, \mathcal{B}(q))} + \frac{w(p)}{2|B|\sum_{q \in P_b} w(q)}$$

Upon receiving a point p, we assign $r(p)$ a uniform random number in the interval $(0, 1)$. This is the threshold for keeping p in our sample S - we keep the point p if and only if $r(p) < Prob(p)$. For each point p, $Prob(p)$ is non-increasing as the stream progresses (this is immediate from the formula and from the fact that clusters never decrease in size since $\mathcal{B}$ is online). Therefore after receiving point p, we update $Prob(s)$ for each $s \in S$ and delete any such s that drop below their threshold: $r(s) \geq Prob(s)$. Once a point crosses the threshold, it may be deleted since $Prob(s)$ is non-increasing and so it will remain below the threshold at all future times. In this way, the construction exactly matches the output as if the offline Algorithm 2 had been used.

Algorithm 4: Input: integer k, ρ-metric space (X, D), stream P of points n from (X, D), online bicriterion $\mathcal{B}$, sample size $m \geq 1$. Output: a coreset S

```
1  S ← ∅
2  for each point p that arrives do
3  |  Assign r(p) a uniform random in (0, 1)
4  |  S ← S ∪ {p}
5  |  Compute Prob(p) according to Line 3 of Algorithm 2
6  |  for each point s ∈ S do
7  |  |  Update Prob(s)
8  |  |  if Prob(s) ≤ r(s) then
9  |  |  |  Delete s from S
10 |  |  Update u(s) according to Line 6 of Algorithm 2
```

Algorithm 3 provides the function $\mathcal{B}$ and a weighted set $M_{i-\lambda}$. Beginning in phase $i - \lambda$, we will begin running Algorithm 4. This outputs the sample S for phases $i - \lambda + 1$ until the current phase i. The following theorem shows that $M_{i-\lambda} \cup S_i$ is a coreset for the stream. Of course, we need to do this construction for each of the $\lambda = O(\log(\frac{1}{\epsilon}))$ most recent phases, so the space gets multiplied by this factor.

We return to definition 6.2 of an r Log-Log Lipschitz function $\tilde{D}$. Given a metric space $(X, dist)$, the space (X, D) where $D = \tilde{D}(dist)$ is a ρ-metric space for $\rho = \max\{2^{r-1}, 1\}$. It is well-known that most M-estimators can be recast as a Log-Log Lipschitz function for a low constant value of r. For example, k-means has $r = 2$.

Theorem 8.5. *There exists a single-pass streaming algorithm requiring* $O(\epsilon^{-O(1)} k \log n(\log n + \log k + \log \frac{1}{\delta}))$ *space that maintains a* (k, ϵ)*-coreset for the* k*-median clustering of an* r*-Log-Log-Lipschitz function* $\tilde{D}$ *on a stream of at most* n *elements with probability at least* $1 - \delta$*.*

Proof. Running Algorithm 2 (producing the M_i) and Algorithm 3 (producing the S_i) in parallel, we will show that $M_{i-\lambda} \cup S_i$ is the desired coreset of the stream.

We already have by Theorem 8.4 that $d_{EM}(M_{i-\lambda}, R_{i-\lambda}) \leq \epsilon \mathsf{OPT}_k(P) \leq \epsilon \bar{g}(P, w, C)$ for any set C of size k. Also, by Theorem 6.6 we have the S_i is an (k, ϵ)-coreset for $P \setminus R_{i-\lambda}$. This is because (in the statement of Theorem 6.6) we are required to carry out the construction of Algorithm 2 using $m \geq \frac{ck(t+\beta)}{\varepsilon^2}\left(d \log t + \log(\beta k) + \log\left(\frac{1}{\delta}\right)\right)$ where $t = \alpha/\psi$. It was shown in Lemma 6.2 that $\psi = (\epsilon/r)^r$. We are using the $(O(1), O(\log(\frac{1}{\epsilon}) \log n))$-approximation $\mathcal{B}$, and in [FL11] it is shown that $d = O(\log n)$ for k-median in a ρ-metric space. So therefore the minimal possible value of m that satisfies the hypotheses of the theorem is $O(\frac{k(\epsilon^{-r} + \log n)}{\varepsilon^2}\left(\log n \log \epsilon^{-r} + \log(k \log n) + \log\left(\frac{1}{\delta}\right)\right))$. Simplifying notation, this is $O(\epsilon^{-O(1)} k \log n(\log n + \log k + \log \frac{1}{\delta}))$.

We write $\mathsf{COST}(A, C)$ to denote $\bar{g}(A, w_A, C)$ to simplify notation; by $A \cup B$ we mean $(A \cup B, w_A + w_B)$. Taking the union, we get that $|\mathsf{COST}(P, C) - \mathsf{COST}(M_{i-\lambda} \cup S_i, C)| \leq |\mathsf{COST}(R_{i-\lambda}, C) - \mathsf{COST}(M_{i-\lambda}, C)| + |\mathsf{COST}(P \setminus R_{i-\lambda}, C) - \mathsf{COST}(S_i, C)|$. The first term is upper-bounded by the Earth-Mover distance, and the second term is upper-bounded by ϵ since S_i is a (k, ϵ)-coreset for $P \setminus R_{i-\lambda}$. So therefore $M_{i-\lambda} \cup S_i$ is a 2ϵ-coreset for the stream P, and the proof is complete after rescaling ϵ. □

The previous theorem has, as special cases, streaming coresets for k-median, k-means, L_p, Cauchy estimators, Tukey estimators, etc. This the first algorithm that does not use merge-and-reduce for building coresets over streams for any of these problems. Moreover, the constant in the exponent for ϵ is small, for the example of k-median the ϵ dependence is $\epsilon^{-3} \log(1/\epsilon)$.

References

[AB99] M. Anthony and P. L. Bartlett. *Neural Network Learning: Theoretical Foundations*. Cambridge University Press, Cambridge, England, 1999.

[AHPV04] P. K. Agarwal, S. Har-Peled, and K. R. Varadarajan. Approximating extent measures of points. *Journal of the ACM*, 51(4):606–635, 2004.

[AMR+12] Marcel R. Ackermann, Marcus Märtens, Christoph Raupach, Kamil Swierkot, Christiane Lammersen, and Christian Sohler. Streamkm++: A clustering algorithm for data streams. *J. Exp. Algorithmics*, 17:2.4:2.1–2.4:2.30, May 2012.

[BF15] Artem Barger and Dan Feldman. k-means for streaming and distributed big sparse data. *SDM'16 and arXiv preprint arXiv:1511.08990*, 2015.

[BLK15] Olivier Bachem, Mario Lucic, and Andreas Krause. Coresets for nonparametric estimation-the case of dp-means. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 209–217, 2015.

[BMO+11] Vladimir Braverman, Adam Meyerson, Rafail Ostrovsky, Alan Roytman, Michael Shindler, and Brian Tagiku. Streaming k-means on well-clusterable data. In *Proceedings of the Twenty-second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '11, pages 26–40. SIAM, 2011.

[BS80] Jon Louis Bentley and James B Saxe. Decomposable searching problems i. static-to-dynamic transformation. *Journal of Algorithms*, 1(4):301–358, 1980.

[CEM+15] Michael B Cohen, Sam Elder, Cameron Musco, Christopher Musco, and Madalina Persu. Dimensionality reduction for k-means clustering and low rank approximation. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*, pages 163–172. ACM, 2015.

[Che08] Ke Chen. A constant factor approximation algorithm for k -median clustering with outliers. In Shang-Hua Teng, editor, *SODA*, pages 826–835. SIAM, 2008.

[Che09a] Ke Chen. On coresets for k-median and k-means clustering in metric and euclidean spaces and their applications. *SIAM Journal on Computing*, 39(3):923–947, 2009.

[Che09b] Ke Chen. On coresets for k-median and k-means clustering in metric and euclidean spaces and their applications. *SIAM J. Comput.*, 39(3):923–947, August 2009.

[COP03] Moses Charikar, Liadan O'Callaghan, and Rina Panigrahy. Better streaming algorithms for clustering problems. In *Proceedings of the Thirty-fifth Annual ACM Symposium on Theory of Computing*, STOC '03, pages 30–39, New York, NY, USA, 2003. ACM.

[DDH+07] Anirban Dasgupta, Petros Drineas, Boulos Harb, Ravi Kumar, and Michael W. Mahoney. Sampling algorithms and coresets for lp regression. *CoRR*, abs/0707.1714, 2007.

[FFK11] Dan Feldman, Matthew Faulkner, and Andreas Krause. Scalable training of mixture models via coresets. In *Advances in Neural Information Processing Systems*, pages 2142–2150, 2011.

[FL11] Dan Feldman and Michael Langberg. A unified framework for approximating and clustering data. In *Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing*, STOC '11, pages 569–578, New York, NY, USA, 2011. ACM.

[FMS07] D. Feldman, M. Monemizadeh, and C. Sohler. A PTAS for k-means clustering based on weak coresets. In *SoCG*, 2007.

[FS05] G. Frahling and C. Sohler. Coresets in dynamic geometric data streams. In *Proc. 37th Annu. ACM Symp. on Theory of Computing (STOC)*, pages 209–217, 2005.

[FS12] Dan Feldman and Leonard J Schulman. Data reduction for weighted and outlier-resistant clustering. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 1343–1354. SIAM, 2012.

[FSS13] Dan Feldman, Melanie Schmidt, and Christian Sohler. Turning big data into tiny data: Constant-size coresets for k-means, pca and projective clustering. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1434–1453. SIAM, 2013.

[GMM+03] Sudipto Guha, Adam Meyerson, Nina Mishra, Rajeev Motwani, and Liadan O'Callaghan. Clustering data streams: Theory and practice. *IEEE Trans. on Knowl. and Data Eng.*, 15(3):515–528, March 2003.

[Guh09] Sudipto Guha. Tight results for clustering and summarizing data streams. In *Proceedings of the 12th International Conference on Database Theory*, ICDT '09, pages 268–275, New York, NY, USA, 2009. ACM.

[HHR11] Frank Hampel, Christian Hennig, and Elvezio Ronchetti. A smoothing principle for the huber and other location m-estimators. *Computational Statistics & Data Analysis*, 55(1):324 – 337, 2011.

[HP11] Sariel Har-Peled. *Geometric approximation algorithms*, volume 173. American mathematical society Providence, 2011.

[HPK07] S. Har-Peled and A. Kushal. Smaller coresets for k-median and k-means clustering. *Discrete Comput. Geom.*, 37(1):3–19, 2007.

[HPM04] S. Har-Peled and S. Mazumdar. On coresets for k-means and k-median clustering. In *STOC*, 2004.

[Hub81] P. J. Huber. Robust statistics. 1981.

[LBK15] Mario Lucic, Olivier Bachem, and Andreas Krause. Strong coresets for hard and soft bregman clustering with applications to exponential family mixtures. *CoRR*, abs/1508.05243, 2015.

[LLS00] Yi Li, Philip M. Long, and Aravind Srinivasan. Improved bounds on the sample complexity of learning. In *SODA*, pages 309–318, 2000.

[LLS01] Y. Li, P. M. Long, and A. Srinivasan. Improved bounds on the sample complexity of learning. *Journal of Computer and System Sciences (JCSS)*, 62, 2001.

[LS10] M. Langberg and L. J. Schulman. Universal ε approximators for integrals. *To appear in proceedings of ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2010.

[Mat89] Jiří Matoušek. Construction of epsilon nets. In *Proceedings of the fifth annual symposium on Computational geometry*, pages 1–10. ACM, 1989.

[PKB14] Dimitris Papailiopoulos, Anastasios Kyrillidis, and Christos Boutsidis. Provable deterministic leverage score sampling. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 997–1006. ACM, 2014.

[SV07] Nariankadu D Shyamalkumar and Kasturi Varadarajan. Efficient subspace approximation algorithms. In *SODA*, volume 7, pages 532–540. Citeseer, 2007.

[Tyl87] David E Tyler. A distribution-free m-estimator of multivariate scatter. *The Annals of Statistics*, pages 234–251, 1987.

[VC71] V. N. Vapnik and A. Y. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory Prob. Appl.*, 16:264–280, 1971.

[Zha11] Z. Zhang. M-estimators. *http://research. microsoft.com/en-us/um/people/zhang/INRIA/ Publis/Tutorial-Estim/node20.html*, [accessed July 2011].